Lead Time Reduction From forecasted to reactive assembly

Master Thesis

Author: Supervisors: W.M. van Dieren

VOORTAAN STEEL GROUP

> University of Twente Dr.ir. M.R.K. Mes Dr.ir. J.M.J. Schutten

Voortman Steel Machinery G.J. Karsten

UNIVERSITY OF TWENTE. Industrial Engineering and Management Production and Logistics Management

Setting the base stock levels for the machine assembly process at Voortman Steel Machinery

Lead Time Reduction

From forecasted to reactive assembly

Setting the base stock levels for the machine assembly process at Voortman Steel Machinery

Date 22 March 2022

Author: W.M. van Dieren s1755617 Industrial Engineering and Management Production and Logistics Management

Supervisory committee: University of Twente Dr.ir. M.R.K. Mes Dr.ir. J.M.J. Schutten

Voortman Steel Machinery G.J. Karsten

University of Twente Faculty of Behavioural, Management and Social sciences Drienerlolaan 5 7522 NB Enschede www.utwente.nl

Voortman Steel Machinery Ozonstraat 1 7463 PK Rijssen +31 (0) 548-536 373 www.voortman.net





Preface

The journey of which this thesis is the endpoint began in 2015 after graduating from a university of applied science in the field of land and water management (i.e., civil engineering). After graduation, I started my first job at one of the 10 largest construction companies in the Netherlands. There I experienced that in reality, things often happen as they do and that optimising processes is a wish that is often expressed but not or only partially achieved. This triggered my interest in why this was the case; I learned and experienced that optimising processes is complex due to parameters and processes influencing each other. Therefore, it turned out that it was not the motivation but the lack of knowledge on how to achieve this. This realisation made me decide to go back to studying and fully committing myself to obtain the knowledge on how to achieve lasting process improvement.

Achieving this knowledge proved to be quite the challenge as I did not pass the last exam of my pre-master in construction management and engineering. Deciding that I still wanted to learn more about optimising processes despite this setback, I did a one-year track to obtain a bachelor in industrial engineering and management. After obtaining this bachelor, I realised that I did not know enough yet to optimise really complex problems. Therefore, I decided to go back to the University of Twente for a pre-master and subsequently this master in industrial engineering and management.

This long process may seem very inefficient for someone who wants to optimise processes. However, it has allowed me to work in different industries and meet many people with different backgrounds, with their own unique skills and methods for solving problems and optimising processes. All these experiences make me better able to link theory and practice, a role I aspire to.

This long 'career' as a student was, of course, not possible without the support of friends and family, who sometimes expressed their doubts if I would ever stop studying but even more often expressed their support and admiration for my determination. Special thanks to all of you at Euros Kano and, in particular, the white water kayakers who made my time in Enschede worthwhile. You made me feel at home and gave me a new lifestyle, so hopefully, I will see many of you on or around the water in the years to come!

Last but not least, I would like to thank my supervisors, without whom this thesis would not have been possible. First of all, I would like to express my gratitude to Gijs Karsten, who made it possible to conduct my graduation project at Voortman Steel Machinery. I am especially grateful for the opportunity to come up with my own research topic and for the flexibility and unwavering support during the process. Next, I would like to express my appreciation to Martijn Mes and Marco Schutten, my supervisors from the University of Twente. I am thankful for their feedback and advice during my research, which definitely has had a positive impact on the quality of this thesis.

And with that, I conclude my preface, and I sincerely hope you enjoy reading about the project I worked on in the previous months.

Wouter van Dieren Enschede, 22 March 2022



Management summary

Voortman Steel Group (VSG) consists of two companies, Voortman Steel Machinery (VSM) and Voortman Steel Construction. VSG has a turnover of 110 million euros and employs more than 500 people, of which 370 work at VSM. VSM focuses exclusively on designing and producing computer numerical controlled (CNC) machines for the structural steel and fabrication industry. The CNC machines are modular so that the customers can configure them to their requirements.

VSM does not want to maintain an anonymous inventory and therefore orders all necessary parts in the required quantity. The downside of this is that the lead times of the parts greatly influence the lead time of the machines. As a result, VSM cannot achieve its objective of completing at least 95% of its machines within 10 weeks. Therefore, VSM decided to forecast expected sales and order parts upfront. Unfortunately, the forecasted sales and machine configurations often deviate considerably from the realised sales and ordered configurations. Consequently, VSM can still not achieve the 10-week lead time objective for at least 95% of its orders.

We hypothesise that pre-assembled machine modules can act as a buffer against the negative impact of part lead times on machine lead times, eliminating the need to forecast expected sales. The storage of parts would also be a viable option to eliminate the negative impact of part lead times on machine lead times. In this study, we opted for modules, as it proved difficult to determine the parts required for the modules accurately. VSM's management is interested in this theory and is thinking of changing the assembly strategy. However, before they do this, they want to have a method that accurately determines which modules should be kept in stock and in what quantity (i.e., base stock levels) while minimising inventory value. Therefore, we formulate the research question as:

How to determine which modules should be kept in stock and in what quantity so that VSM can guarantee a 10-week lead time for at least 95% of the orders while minimising inventory value?

We started by studying the literature to determine which aspects are normally considered by assembly strategies for modular products and what methods are already developed to determine adequate base stock levels. From the literature, we learned that VSM applies the assemble-to-order (ATO) strategy and that product design, assembly layout design, and the applied planning method need to be covered adequately to ensure the successful realisation of a flexible ATO. Determining the most suitable product modules requires in-depth technical knowledge of the product, which is difficult to obtain quickly. Therefore we decided to use the machine modules currently distinguished by VSM. The fixed position assembly layout that VSM currently applies proved to be ideal for the relatively low production volumes with long setup times of the VSM machines. With VSM's management, we decided that a priority rule-based scheduling heuristic best represents VSM's scheduling process.

Several methods have been developed for determining base stock levels; however, the common consensus of the developers of these models is that determining optimal base stock levels is difficult and computationally tedious, especially when considering component commonality. Due to the complexity of this problem, most developed methods are highly theoretical and simplified. Therefore we decided to develop a practical method to determine base stock levels for a multi-product ATO system with complex machine configurations in order to be able to determine the base stock levels for VSM.

In order to develop such a method, we had to decide on the solution approach. We did not consider enumeration suitable or even possible due to the enormous solution space. Also, exact optimisation methods are not applicable as Bienstock and Özbay (2008) proved that determining robust base stock

levels is NP-hard. Therefore we decided that heuristic approaches were the most viable option for generating robust sets of base stock levels. We came up with two heuristic approaches that we deemed to be the most feasible.

The first approach is a genetic algorithm (GA); we think that GAs might be suitable for determining robust sets of base stock levels because of their ability to deal with real-life size problems and their ability to use historical data to guide the search to the best performing region within the solution space (Daniel and Rajendran, 2005). The second approach is a combination of a GA and local search algorithm; we refer to this approach as the local search approach. We opted for this combination as GAs are not well suited for fine-tuning solutions (i.e., local search), which are very close to optimal ones (Martinez and Lozano, 2007). The GA guides the search to the best performing region within the solution space, after which the local search takes over to analyse this region. We selected local search as it has proven to be very successful in determining near-optimal and sometimes even optimal solutions for difficult real-life problems with enormous solution spaces (Aarts and Lenstra, 2003; Dumitrescu and Stützle, 2003).

Both approaches are heuristics and not exact algorithms; therefore, they cannot guarantee that they found the optimal set of base stock levels; instead, they can generate sets of base stock levels capable of meeting the objective. The objective is to complete at least 95% of the machines of future demand scenarios within 10 weeks. Therefore, we let both approaches generate multiple sets of base stock levels and select the most cost-effective set that can meet the objective as the best solution. To ensure that the selected most cost-effective set of base stock levels is a good solution, we need to gather at least 25,000 possible solutions. The number of 25,000 is based on the number of different module types we need to take into account and the dual objective of minimising the inventory value while meeting the lead time objective.

Conducting experiments with VSM's real assembly system is not practical; therefore, we had to find another method to determine the performance of the sets of base stock levels generated by both approaches. Together with VSM's management, we decided that using a simulation model of VSM's purchase, inventory management, and assembly process would be the most suitable for determining what machine lead times can be realised for a given demand scenario and set of base stock levels.

With the help of the simulation model, we analysed the performance of the developed GA and local search approach. We analysed if the most cost-effective set of base stock levels generated by both approaches were capable of ensuring with 95% certainty that at least 95% of the machines of future demand scenarios can be completed within 10 weeks. We did this by simulating multiple future demand scenarios with the sets of base stock levels to see how they performed; by conducting statistical tests (i.e., one-sample t-tests), we determined that both approaches can create sets of base stock levels that can fulfil the lead time objective. However, besides meeting the required lead-time objective, we also want to minimise the required inventory value. Therefore, we analysed 1,000 sets of base stock levels generated by each approach. Due to time constraints, we analysed 1,000 sets of base stock levels and not 25,000 as creating them for both approaches would take months. Based on a sample size analysis, we concluded that a sample consisting of 1,000 sets of base stock level is large enough for the analysis of both approaches. We plotted the 1,000 sets of base stock levels generated by both approaches against each other in Figure 1, sorted in ascending order. The inventory value of the sets of base stock levels can be read from the y-axis, and the number of the sets are stated on the x-axis. With the 1,000 sets of base stock levels of both approaches, we conducted a statistical test (i.e., two-sample t-test) and concluded that the local search approach creates significantly cheaper sets of base stock levels.





Figure 1: The inventory values in ascending order of the 1,000 sets of base stock levels generated by the GA and local search approach.

In practice, it is common to apply an equal fill rate approach due to the lack of practical approaches on setting individual stock-keeping unit (SKU) levels in a multi-SKU environment. The equal fill rate approach sets the individual item fill rates equal to the targeted fill rate of the entire inventory. For the situation under study at VSM, this would mean that we set the individual fill rates of all the 288 module types we take into account to 95%. The downside of the equal fill rate approach is that this method is inaccurate and leads to higher inventory values than necessary because not all SKUs have the same impact on the systems fulfilment performance, i.e. some SKUs are overstocked (Teunter et al., 2017). To set the inventory value of the sets of base stock levels generated by the GA and local search approach into perspective, we compare them with the inventory value of the set of base stock levels generated by the equal fill rate approach. The inventory value of the set of base stock levels generated by the equal fill rate approach is 48.3 million euros. Based on this, we conclude that the set of base stock levels generated by the equal fill rate approach is outperformed in terms of cost-effectiveness by all sets of base stock levels generated by the local search approach and most of the sets of base stock levels generated by the GA approach.

To get an overview, we plotted the most cost-effective sets of base stock levels generated by the GA and local search approach against the set of base stock levels generated with the equal fill rate approach in Figure 2. From this figure, we conclude that the local search approach is the best approach for determining cost-effective, robust base stock levels. The GA performs a bit worse but is still applicable. However, the equal fill rate approach we do not deem suitable for determining cost-effective base stock levels.



Figure 2: Inventory value of most cost-effective sets of base stock levels generated by the GA and local search approach and the set of base stock levels generated with the equal fill rate approach.

The last step was to decide if the approaches we developed and tested with the situation under study at VSM really contribute to the literature and help fill the research gap. Based on the assessment of both approaches, we concluded that they are both suitable for generating robust sets of base stock levels for full-sized real-world problems but that the practical implementability of both approaches could be better. Nonetheless, we consider both approaches a useful contribution to the literature.



Table of Contents

Prefaceii
Management summaryiii
List of figures and tablesx
Figuresx
Tablesxi
Glossaryxii
1. Introduction
1.1 Company introduction1
1.2 Problem context
1.3 Problem statement & Research objective
1.4 Research scope
1.5 Research questions 4
1.5.1 Question 1 4
1.5.2 Question 2 4
1.5.3 Question 3
1.5.4 Question 4
1.6 Applied solution approaches5
2. Current situation
2.1 Product range
2.1.1 Multi-Systems Integration
2.1.2 Transition from ETO to CTO7
2.2 Facility layout
2.3 Involved departments
2.3.1 Sales department9
2.3.2 Operations department10
2.3.3 Purchase department 10
2.4 Assembly and purchase planning procedures11
2.5 Forecast accuracy
2.6 Conclusion
3. Literature review
3.1 Aspects to consider when determining base stock levels15
3.1.1 Assemble to order
3.1.2 Modular product design
3.1.3 Assembly layout design

	3.1.4 Planning methods	. 19
	3.2 Approaches for determining base stock levels	. 21
	3.3 Research gap	. 23
	3.4 Solution approaches	. 23
	3.5 Conclusion	. 24
4.	Solution design	. 27
	4.1 Problem definition	. 27
	4.1.1 Assembly and test cells	. 28
	4.1.2 Assembly and test hours	. 28
	4.1.3 Base stock levels	. 29
	4.2 Solution approach	. 30
	4.3 Synergy between heuristic approaches and the simulation model	. 31
	4.3.1 Synergy GA and simulation model	. 31
	4.3.2 Synergy local search and simulation model	. 33
	4.4 Heuristic approaches	. 34
	4.4.1 Genetic algorithm (GA) approach	. 34
	4.4.2 Local search approach	. 37
	4.5 Simulation model	. 39
	4.5.1 Scenario generation	. 39
	4.5.2 Scheduling algorithm	. 40
	4.5.3 Validation	. 43
	4.6 Conclusion	. 44
5.	Analysis	. 46
	5.1 Procedure for generating sets of base stock levels	. 46
	5.2 Data collection	. 47
	5.3 Performance analysis solution approaches	. 48
	5.3.1 Lead times	. 49
	5.3.2 Inventory value	. 51
	5.4 Fulfilment of research gap	. 56
	5.5 Conclusion	. 57
6.	Conclusions and discussion	. 59
	6.1 Conclusion	. 59
	6.2 Recommendations	. 62
	6.3 Limitations	. 63
	6.4 Future research	. 64



References	65
Appendix	69
A. VSM's product range	70
B. Histogram forecasting accuracy	88
C. Modules of the considered VSM machine types	88
D. Poisson distribution and seasonality per machine type	88
E. Sample size analysis	88
F. Base stock levels of the best solutions	90

List of figures and tables

Figures	
Figure 1.1 - Subsidiaries and dealer network of VSM (Voortman Steel Group, 2021).	1
Figure 1.2 - Example per machine segment (Voortman Steel Group, 2021).	2
Figure 2.1 - Machining processes per product segment.	6
Figure 2.2 - Multi-System Integration (MSI).	7
Figure 2.3 - ETO and CTO process.	8
Figure 2.4 - VSM assembly facility with distinguished assembly areas.	9
Figure 2.5 - In the DNA tool, stored project data.	11
Figure 2.6 - Part of the assembly planning of the Operations department.	13
Figure 2.7 - Workload graph from week 17 onwards, left: before replanning right: after replanning. Week 27, 28 and 29 is the summer holiday period.	14
Figure 2.8 - Example of accuracy determination for machine A.	14
Figure 2.9 - The relevant part of the forecasting accuracy histogram, see appendix C for the total histogram.	15
Figure 3.1 - Distinguished CODP (Olhager 2010).	17
Figure 3.2 - Hierarchical planning-and-control frameworks by Hans et al. (2007).	21
Figure 3.3 - System fulfilment performance VS Inventory value based on (Cardós et al., 2013).	22
Figure 3.4 - left, Local neighbourhood search converging to local optima. Right, local(x_1) and global(x_2) optimum for search space f(x). (Bani-Hani, 2020)	25
Figure 4.1 - Configure to order options V807.	28
Figure 4.2 - Flowchart of the collaboration between the GA approach and the simulation model.	33
Figure 4.3 - The required minimum number of simulation trials per amount of decision variables (Thengvall, 2019).	34
Figure 4.4 - Flowchart of the collaboration between the local search approach and the simulation model.	34
Figure 4.5 - Applied mating probability function.	36
Figure 4.6 - Schematic representation of the chromosome-like data structures of the individuals in a generation.	36
Figure 4.7 - k-point crossover, k can be any non-negative integer.	36
Figure 4.8 - left: flow chart of the GA we designed to create an initial set of base stock levels. right: flow chart of the GA we designed to improve on an insufficient set of base stock levels.	38
Figure 4.9 - Flow chart of the Local search algorithm we applied to improve insufficient sets of base stock levels.	40
Figure 4.10 - Process for generating stochastic demand scenarios.	41
Figure 4.11 - Schematic view of priority rule-based heuristic with explanations of the different sets.	42
Figure 4.12 - Triangular probability distribution and formulas (Law, 2015).	43
Figure 4.13 - Gannt chart of a one-year assembly schedule for VSM generated by the priority rule-based scheduling heuristic.	43
Figure 4.14 - The priority rule-based scheduling heuristic we designed to create assembly schedules representative of VSM's scheduling process	44



Figure 5.1 - The best (i.e., cheapest) inventory value per number of gen for both the GA and the local search approach.	erated sets of base stock levels	49
Figure 5.2 - Graphic statistical summary report of the 250 executed sim (i.e., scheduled one-year demand scenarios) for both appro	ulation runs aches' best set of base stock levels.	50
Figure 5.3 - Report of the one-sample t-test for the most cost-effective	solution of both approaches.	51
Figure 5.4 - Report of the two-sample t-test.		52
Figure 5.5 - The inventory values in ascending order of the 1,000 sets of by the GA and local search approach.	f base stock levels generated	53
Figure 5.6 - The number of modules per set of base stock levels in Figur	e 5.5.	54
Figure 5.7 - Graphic statistical summary report of the inventory value o sets of base stock levels per approach.	f 1,000 generated	54
Figure 5.8 - Report of the two-sample t-test.		55
Figure 5.9 - Graphic statistical summary report of the 250 executed sim	ulation runs for the equal fill rate approach.	56
Figure 5.10 - Report of the one-sample t-test.		56
Figure 5.11 - Inventory value of most cost-effective sets of base stock le approach and the set of base stock levels generated with t	evels generated by the GA and local search he equal fill rate approach.	57
Figure 6.1 - Inventory value of most cost-effective sets of base stock lev approach and the set of base stock levels generated with th	vels generated by the GA and local search e equal fill rate approach.	62
Tables		
Table 2.1 - Sales results of the last four years.		7
Table 4.1 - Amount of test and assembly cells per zone and machine type	pe.	29
Table 4.2 - The amount of testing and assembly hours required per mac and the weeks required for assembly and testing.	chine type	30

Glossary	
ATO	Assemble To Order
BOM	Bill Of Materials
CI	Confidence Interval
CNC	Computer Numerical Controlled
CODP	Customer Order Decoupling Point
СТО	Configure To Order
DOQ	Discrete Order Quantity
ERP	Enterprise Resource Program
ETO	Engineer To Order
GA	Genetic Algorithm
JIT	Just In Time
LS	Local Search
MPD	Modular Product Design
MPSM	Managerial Problem Solving Method
MSI	Multi Systems Integration
MTO	Make To Order
MTS	Make To Stock
NP	Non-Polynomial
PBA	Population Based Algorithms
RMS	Reconfigurable Modular Systems
SKU	Stock Keeping Unit
SSBA	Single Solution Based Algorithms
VSG	Voortman Steel Group
VSM	Voortman Steel Machinery



1. Introduction

In this research, we develop a method to determine which modules need to be kept in stock and in what quantity at Voortman Steel Machinery (VSM). Section 1.1 provides a company introduction, followed by the problem context in Section 1.2. Section 1.3 discusses the problem statement and research objective. Section 1.4 defines the research scope, and Section 1.5 states the formulated research questions.

1.1 Company introduction

Voortman Steel Group (VSG) consists of 2 companies, VSM and Voortman Steel Construction. VSG has a turnover of 110 million euros and employs more than 500 people, of which 370 work at VSM.

Voortman was founded in 1968 in Rijssen in the Netherlands by the brothers Voortman. For the first two years, they had a broad focus on all kinds of machinery, but in 1970 the focus shifted towards mechanisation, which led to the company's rapid growth. After six years, the company focus broadened, and VSG started designing and building steel construction frames. Four years later, in 1980, the family Voortman split the company between Voortman Automatisering B.V., which is now known as Voortman Steel Machinery B.V and Voortman Staalbouw B.V., currently known as Voortman Steel Construction B.V.

Since 1995, VSM has focused exclusively on designing and producing computer numerical controlled (CNC) machines for the structural steel and fabrication industry. This specialisation resulted in impressive growth and enabled VSM to expand internationally, e.g., opening subsidiaries in Germany, the United Kingdom, the USA, Russia, France and Poland. Furthermore, VSM set up a dealer network spanning six continents (see Figure 1.1. Manufacturing of all the machines takes place at VSM's headquarters in Rijssen. Subsidiaries are engaged in sales and service. Currently, 35% of VSM's sales activities occur in Europe and 65% outside Europe (Voortman Steel Group, 2021).



Figure 1.1: Subsidiaries and dealer network of VSM (Voortman Steel Group, 2021).

In 2011 VSM widened its scope by acquiring Maschinenfabrik Bach, which specialised in plate processing CNC machines, and integrated it into VSM, thereby expanding its machinery range to:

- Beam processing;
- Flat & angle processing;
- Surface treatment;
- Plate processing.

See Figure 1.2 for an example per machine segment.



Figure 1.2: Example per machine segment (Voortman Steel Group, 2021).

1.2 Problem context

The machines of VSM have a modular design, i.e. the machines are composed of modules, and the modules are composed of parts. The use of a modular machine design enables to configure machines to order with predesigned modules. Since the machines can have a wide variety of configurations, and due to their relatively low manufacturing volumes and high value, production to stock is not suitable.

VSM does not keep free to use (i.e., anonymous) part inventory as it applies a lot-for-lot inventory strategy, meaning that all the necessary parts for a machine are ordered based on a customer order. By applying this strategy, VSM enables Just In Time (JIT) delivery of parts to keep inventory to a minimum. However, this strategy also has a major drawback for VSM, as individual planning is necessary to ensure that each part arrives on time.

Since the inventory is linked to specific customer orders, all purchase activities start as a result of a machine order, which is called a cold start at VSM. Due to this cold start, the lead time of the required parts greatly influences the lead time of the respective modules and thus the machine's lead time. VSM would like to guarantee a lead time of 10 weeks or less in 95% of the time to increase customer satisfaction. VSM defines lead time as the time between a customer order and the machine being ready for transport. Currently, no machines are finished within 10 weeks. To counter these long lead times, VSM started forecasting expected sales and machine configurations in order to be able to plan ahead.



This forecasting procedure enables VSM to order parts before a customer orders a machine. However, due to the wide variety of machine configurations, it is not easy to make accurate predictions. Therefore, it regularly happens that the predicted sales and machine configurations deviate considerably from the realised sales and ordered configurations. As a result, the assembly plan needs to be adjusted frequently to prioritise machines that customers have ordered. Furthermore, the Purchase department often needs to send high-priority purchase requests for the parts not yet ordered if an ordered configuration.

The frequent replanning of the assembly activities means that the purchase plan also must be adjusted frequently. To cope with the frequent replanning, the Purchase department decided not to apply the implemented JIT strategy because the JIT strategy orders all parts separately, which is time-consuming, especially with frequent changes. Therefore, the Purchase department chose to combine orders and order everything in one large order per supplier to save time. The delivery of the combined orders must occur before the first part in the order is needed, which results in a large stock. This approach, therefore, contradicts the arguments for choosing the lot-for-lot strategy in the first place, as there is now a large inventory, and the parts are no longer delivered JIT.

1.3 Problem statement & Research objective

We deduct the following problem statement from the problem context in Section 1.2:

The current assembly strategy of VSM leads to frequent revision of the assembly plan, which causes disruptions in the supply chain, making it difficult to ensure a lead time of no more than 10 weeks for 95% of the orders.

We have the theory that pre-assembled machine modulesⁱ can act as a buffer against the negative impact of part lead times on machine lead times, eliminating the need to forecast expected sales (i.e., orders are not planned until they are actually placed). Therefore, VSM's management is thinking of changing the assembly strategy by making the strategic decision to keep pre-assembled machine modules in stock. However, before they do, they would like to have a method capable of accurately determining which modules should be kept in stock and in what quantity (i.e., base stock levels). Therefore, we formulate the research objective as:

Develop a method for VSM to determine which modules should be kept in stock and in what quantity, so that a 10-week lead time can be guaranteed for at least 95% of the orders, while minimising inventory value.

1.4 Research scope

We only include the departments directly involved in the assembly process in this research, limiting the scope of this research to the Sales, Operations and Purchase departments. The Sales department sells the machines, the Operations department organises the assembly of the machines, and the Purchase department sources all the necessary parts for the machines.

Machine types that VSM is phasing out and machine types that VSM did not sell in the past year are excluded from this research by VSMs management. In Section 2.1, which covers VSM's product range, we determine to which machine types this applies. We deem the lead times of the parts needed for a machine module equal to the part within the module with the longest lead time. VSM's management

ⁱ The storage of parts would also be a viable option to eliminate the negative impact of part lead times on machine lead times. In this research, we opted for modules, as it proved difficult to accurately determine the parts required for the modules.

also has indicated that there is no plan in the near future to change the assembly layout, so we consider the current assembly layout a given.

1.5 Research questions

To accomplish the research objective, we formulate the following main research question:

How to determine which modules should be kept in stock and in what quantity so that VSM can guarantee a 10-week lead time for at least 95% of the orders while minimising inventory value?

To answer the main research question and fulfil the research objective, we apply the managerial problem-solving method (MPSM) of Heerkens and van Winden (2012). We divide the solution process of the MPSM into 5 stages being: (1) current situation, (2) literature review, (3) solution design, (4) analysis of results and (5) conclusion and recommendations. The first four stages all represent one of the research questions discussed in the Subsections below. Stage 5 we cover in Chapter 6, where we conclude this research and present recommendations for further research.

1.5.1 Question 1

Question 1 has the purpose of providing insight into the current situation at VSM regarding the assembly process. To this end, we first analyse the product range of VSM and determine which machines are included in this research. After that, we analyse the assembly layout and the assembly process from forecasting to assembly. We end with determining the forecasting accuracy; Chapter 2 answers Question 1 and its sub-questions.

Question 1 What is the current situation at VSM regarding the machine assembly process?

- a. What is the product range of VSM?
- b. What is the assembly layout of VSM's assembly facility?
- c. What are the steps of VSM's assembly process?
- d. What is the accuracy of the forecasting procedure?

1.5.2 Question 2

Question 2 aims to acquire a theoretical background by conducting a literature review. To achieve this theoretical background, we start by analysing assembly strategies for modular products to determine which aspects are important for a well-functioning assembly strategy of modular products. We do this because we want to consider these aspects when determining which modules should be kept in stock and in what quantity. After that, we study the literature to determine which methods are suitable for determining modular product base stock levels. Chapter 3 answers Question 2 and its sub-questions.

Question 2 What relevant knowledge from the academic literature can we use to determine which modules and in what quantity VSM needs to keep in stock?

- a. What aspects are normally considered by assembly strategies for modular products?
- b. Which methods are suitable for determining the base stock levels of modular products?



1.5.3 Question 3

Question 3 aims to develop a method capable of determining which modules should be kept in stock and in what quantity so that VSM can guarantee a 10-week lead time for at least 95% of the orders while minimising inventory value. To this end, we first specify the problem definition in order to determine what aspects we take into account and how. After that, we determine which solution approaches we think are the most suitable for determining what modules should be kept in stock and in what quantity to reach the objective; Chapter 4 answers Question 3 and its sub-questions.

Question 3 How to determine which modules need to be kept in stock and in what quantity at VSM?

- a. What aspects do we take into account when determining which modules to stock?
- b. Which solution approaches are the most suitable for determining the required base stock levels?

1.5.4 Question 4

Question 4 aims to analyse the approaches designed in Question 3 for determining which modules should be kept in stock and in what quantity so that VSM can guarantee a 10-week lead time for at least 95% of the orders while minimising inventory value. To this end, we first need to decide how to determine and compare the performance of the different approaches. After that, we determine the performance of the different approaches and analyse which performs best and if it is suitable for VSM. Chapter 5 answers Question 4 and its sub-questions.

Question 4 What is the best approach for determining base stock levels at VSM?

- a. How to determine and compare the performance of the different approaches?
- b. How do the different approaches perform?

1.6 Applied solution approaches

We apply two approaches for the generation of sets of base stock levels that should be able to ensure that VSM with 95% certainty can complete at least 95% of the machines of future demand scenarios within 10 weeks. The two approaches we apply are a genetic algorithm (GA) and a local search-based procedure.

Testing the generated sets of base stock levels with VSM's real assembly system is not practical due to disruptions, costs, and time durations. Therefore, we apply a simulation model which simulates VSM's purchase, inventory management and assembly process. The simulation model is as realistic as possible in order to function as a digital twin of the real system. The performance of the sets of base stock levels given as input to the simulation model is measured in machine lead times. We store a set of base stock levels as a possible solution when the percentage of machines completed within 10 weeks is at least 95%. Otherwise, we return the set of base stock levels to the applied approach for improvement. The approaches and the simulation model are discussed in more detail in Chapter 4 about the solution design.

2. Current situation

In Chapter 1, we gave an introduction to this research. In this chapter, we analyse the current situation at VSM. Thereby answering the first research question: 'What is the current situation at VSM?'. The chapter starts with identifying and analysing the product range of VSM in Section 2.1. Section 2.2 focuses on the layout of the assembly facility, and Section 2.3 analyses the departments involved assembly process. Section 2.4 discusses the planning procedure of the Operations and Purchase departments, while Section 2.5 focuses on forecasting accuracy. Section 2.6 concludes this chapter.

2.1 Product range

VSM produces a complete range of efficient and rigid machines for the structural steel and fabrication industry. VSM divides their machines into four product segments: beam processing, plate processing, flat & angle processing and surface treatment. Figure 2.1 states all the machining processes that the machines can perform per product segment. Appendix A indicates per machine what machining operations it can perform as this is machine specific.



Figure 2.1: Machining processes per product segment.

To get a feeling for the manufacturing volumes of VSM, we present the sales result from the last four years per product segment in Table 2.1. In this table is indicated with grew which machines types we exclude for this research. Excluded are the Drill V630-1250, Drill/Saw V630-1250, V704, V808, V200, V302 and V70 due to being phased out. Also excluded are the V613-1050, Drill V631-1250, Drill/Saw V631-1250, V505M, V505T and VP on the instructions of VSM's management due to either no sales in the past year or insufficient sales volumes to take them into account. Also not taken into account for other reasons are the V303 as this machine is new and expected sales are hard to predict and the V71 and V3100 as the production of these machines is mainly outsourced.

Table 2.1: Sales results of the last four years.

Confidential

2.1.1 Multi-Systems Integration

Sometimes the sales agents of VSM sell stand-alone machines, but often they sell complete production lines called Multi-Systems Integration (MSI). These MSIs consist of multiple machines with the



necessary roller conveyors, cross transports, and product buffers to enable fully automated production. VSM's VACAM software facilitates automated production by centrally controlling the production line from one central computer. Figure 2.2 shows a possible MSI configuration. This configuration consists of an infeed and roller conveyor, which transports steel beams to the V630 beam processing machine (left), which marks, drills and mils the beams. Then the beams are transported to the VB1050 saw (middle), where they are cut-to-length. Then the beams are transported to the outfeed right after the saw if they do not require additional processing. Otherwise, they are transported to the V807 coping machine (right) and eventually to the outfeed after the coping machine.



Figure 2.2: Multi-System Integration (MSI).

VSM considers roller conveyors, cross transports, and product buffers as support equipment and not as machines. Therefore, we do not take them into account in this research.

2.1.2 Transition from ETO to CTO

In the past, VSM engineered-to-order (ETO) its machines to ensure the fulfilment of customer requirements. However, ETO is a time-consuming and expensive process as it starts from scratch and requires skilled mechanics and engineers. Another disadvantage of ETO is that the customer order decoupling point (CODP) is upstream, which leads to long lead times. VSM, therefore, decided to opt for an assemble-to-order (ATO) strategy and convert its machines to a modular design with standard modules that enable a configure-to-order (CTO) approach, which moves the customer order decoupling point more downstream. Figure 2.3 shows the differences between the ETO and CTO processes. The designed standard modules enable the independent creation, modification, replacement or exchange of modules to change the configuration of a machine.

Sometimes, however, despite the introduction of CTO, the sales agents still sell machines with impossible configurations. The reason for this is the high degree of complexity of the MSI production lines, which require much in-depth technical knowledge of the machines and support equipment. VSM, therefore, designed a tool for the sales agents, called 'DNA', which enables them to design MSI lines that consist of only standard predesigned machine modules.



Figure 2.3: ETO and CTO process.

2.2 Facility layout

Assembly layouts have a major impact on production efficiency, and thus the required inventory levels, as the arrangement of machines and other equipment greatly influences product lead times and needed intermediate materials (i.e., modules) (Gayam et al., 2020).

VSM uses a fixed-position layout for the assembly of the machines. Fixed-position layouts are applied when the project is considered immobile during manufacturing. The project remains in the same location and workers and equipment come to the project, which is common in heavy industry (e.g., ships, aeroplanes, buildings). The benefits of fixed position layouts are adaptability of project configuration and minimised project and component handling. Downsides are the large number of square meters needed, low production volumes, and complex space and activity scheduling processes (Gaither and Frazier, 2004).

VSM has chosen a fixed position assembly layout as their machines are heavy and unwieldy while the manufacturing volumes are relatively low. The fixed-position layout enables VSM to produce at a pace the market requires while keeping inventory low. VSM divided its assembly facility into separate fixed position assembly areas with the capacity for multiple machines, see Figure 2.4. These fixed position assembly areas focus on improving efficiency, quality and throughput by bringing together all the required equipment, people, and processes needed to manufacture one or more machine types. VSM based the segmentation of the manufacturing facility upon historical demand, sales forecasts, assembly techniques, materials involved, and available resources.

VSM distinguishes 'make' and 'buy' parts. Make parts are custom made parts, outsourced to third parties or internally manufactured by the Parts Manufacturing department of VSM. VSM orders the buy parts from the catalogues of external suppliers. The make parts are stored on the warehouse floor in racks, and the buy parts are stored in either the automated storage and retrieval pallet system or one of the vertical lift modules.

A section of the assembly facility is reserved for the Research and Development department to design and test new machines. The remaining space in the manufacturing facility is for the storage of assembled modules or machines.





Figure 2.4: VSM assembly facility with distinguished assembly areas.

2.3 Involved departments

The departments involved in the assembly of the VSM machines are the Sales, Operations and Purchase departments. The Sales department sells the machines, the Operations department generates the assembly plans and schedules, and the Purchase department sources all the necessary parts for the scheduled machines. The working methods applied by the departments involved in the assembly of the VSM machines are interdependent and exert influence on each other. Therefore they are analysed in more depth in the subsections below.

2.3.1 Sales department

A customer requires a stand-alone machine or a production line to fulfil specific machining processes (e.g., sawing drilling, milling). Based on these requirements, a sales agent from VSM configures a machine or an MSI with the help of the 'DNA' tool. The configurator enables sales agents to configure customised machines and MSI's consisting out of standard modules. The configurator is basically a list of machine specific questions (e.g., installation country, alignment with production line, control options, weight and lengths of processed products). However, sometimes the sales agent finds it necessary to configure a machine that deviates from the standard configuration options to realise a deal.

The machines of VSM are costly, and therefore customers often take their time when considering which machines to order and in which configuration. The sales agents keep a record of the companies considering buying a machine or MSI. The sales agent also indicates what they think the probability is that a deal will occur in this record. Based on these probabilities, the Operation department may

decide to already plan a machine that is not yet sold to reduce the machine's lead time. Section 2.4 discusses this process in more detail.

2.3.2 Operations department

The tasks of the Operations department are threefold. The first one is updating the rolling forecast assembly plan, which the Operations department manager does together with the manager from the Sales department during forecast meetings. Section 2.4 discusses this planning process in more detail. The second task is transforming the rolling forecast assembly plan into an operational assembly schedule that indicates which machines are assembled in what assembly cell and by whom. The third task is the preparation of the on-site installation of the machines.

Machines can be added to the assembly plan before a customer actually orders a machine; Section 2.4 discusses this in more detail. However, tasks two and three only start after the required advance payment for a machine has been received. After receiving the required payment, the Operations department adds the machine to the assembly schedule and prepares the required documents and models based on the project data stored in the DNA tool; see Figure 2.5.

Although the machines are composed out of standard CTO modules, on-site installation still requires some customer-specific engineering. For example, the machine's software configuration and calibration must be prepared to be operational immediately after installation. A simulation model is applied to enable this to be partially done from the office. Besides that, connecting the machines to the local electricity grid requires specific electrical switch boxes and cables. The required cable lengths are determined based on the AutoCAD drawings of the project layout. The last things retrieved from the DNA tool are documents regarding the assembly, such as the promised installation date and the bill of materials (BOM). The Operations department set the materials inside the BOM as requested in the used enterprise resource program (ERP) called SAP. The Purchase department is responsible for sourcing the requested parts.



Figure 2.5: In the DNA tool, stored project data.

2.3.3 Purchase department

VSM purchase department currently applies a discrete order quantity (DOQ) inventory management method, commonly called lot-for-lot. DOQ means that products are produced, and parts are purchased



in the exact amount needed to minimise inventory value (Stevenson, 2007). DOQ is commonly utilised in the case of expensive items with intermitted demand. The downside of DOQ is that because the objective is to minimise inventory value, it maximises the ordering costs by increasing the number of orders placed. Therefore, DOQ only comes to its right in case of neglectable ordering costs and replenishment times (Gosrani and Kolekar, 2017).

VSM applies the DOQ inventory management method as it does not want to keep free to use anonymous inventory. The reason for this is the wide variety of demand patterns and configurations that the VSM machines can have. Therefore keeping anonymous inventory would result in high inventory value as the most expensive parts of the VSM machines are often machine specific. Unfortunately for VSM, the replenishment times of the parts are not neglectable as almost all modules have one or more parts with a replenishment time of at least 6 weeks. Therefore DOQ is not ideally suited for VSM as now the lead time of the required parts greatly influence the lead time of the respective modules and thus the lead times of the machines. To overcome this, VSM started forecasting expected sales and machine configurations; Section 2.4 discusses this in more detail.

However, the unfortunate reality is that the currently applied forecasting method and connected assembly planning method do not work well with this DOQ inventory strategy due to the frequent replanning of the assembly plan. The replanning causes problems as parts are ordered and delivered per assembly slot; thus, the replanning of an assembly slot leads to the replanning of all parts needed for that assembly slot. Another problem is that suppliers often deliver parts for several machine types. Therefore, it often happens that the necessary parts for several machines are ordered simultaneously from one supplier with different delivery times. A supplier books these ordered parts under one order number but with different delivery times. In case of rearrangement of one or more assembly slots in the assembly planning, there is a chance that some parts in already outstanding orders need a different delivery time, while others do not. If this is the case, intensive communication between the Purchase department and the suppliers is essential, which is frustrating and time-consuming.

Another problem is sometimes ordered parts are not needed anymore due to changed machine configurations by the customers. However, these parts are often still delivered as cancelling the orders is not always possible. The reason for this is that a lot of parts are specially made for VSM, meaning that when the production of these has started or preparations have been made, VSM is obliged to take them. When parts are not custom-made for VSM, orders are still not always cancelled as the Purchasing department is too busy and stores them for future machines. The educated guess of VSM's management is that there is for approximately 7 million euros of parts in storage. This storage often does not contribute to the machines' shorter lead times as there are often still parts missing, which indicates the need to store the right parts.

2.4 Assembly and purchase planning procedures

VSM currently applies a rolling forecast planning method, which enables VSM to plan continuously (i.e., forecast) over a set time horizon. Figure 2.6 shows an example of the rolling forecast assembly plan for various machine types—the assembly plan groups the machines by the assembly area used. Per machine type, assembly slots are created based on the number of assembly cells. These slots indicate the week in which the machine should be finished. They do not cover the total assembly period. The created assembly slots are indicated in the assembly planning by the colour blue.

During the forecast meetings, the managers of the Sales and Operations departments release assembly slots per machine type based on expected sales. An assembly slot may be released if the sales agent concerned believes that the probability of the customer placing the order is greater than 50%. The

sales agent also indicates which machine configuration they think the customer is likely to order. The managers then discuss whether or not to release an assembly slot. If an assembly slot is released, the Operations department sets the parts for the predicted configuration as requested in SAP, and the Purchase department orders them. Released assembly slots are indicated in the assembly planning by the colour green. Released assembly slots can also be used for other customers for whom the managers had decided not to release an assembly slot yet because they were not yet expected their order.

The method of ordering parts before a machine is actually ordered ensures that the lead time of the machine is shorter than the 'cold start' period. The 'cold start' period is the time required to produce a machine when starting from scratch, e.g., ordering parts, assembling the machine, testing the machine and preparing it for transport.

When a customer places an order, the assembly slot to which the customer order is allocated is changed from green to red and receives a week number indicating which week the machine must be ready for transport to be delivered on time. If there are no green assembly slots available, additional assembly slots are released as orders are always accepted.

A machine configuration is considered final after the required advance payment from a customer has been received. After receiving the payment, the Operations department adds the machine to the operational assembly schedule, and not yet ordered parts are ordered with urgency by the Purchase department. The lead times of the additionally ordered parts can have a negative effect on the eventual lead time of the machine.



Figure 2.6: Part of the assembly planning of the Operations department.

Customers sometimes decide not to place an order at VSM after all or hesitate and postpone their order. Due to this, VSM releases more assembly slots than they have assembly FTEs available, hoping



this balances out in the end. Unfortunately, this is often not the case. Therefore replanning is frequently necessary to reduce the workload in specific weeks. The method used to determine the workload calculates the sum-product of the number of machines and the required assembly time per machine type. Another reason for replanning is that the minimum time required for engineering and ordering customer-specific parts for on-site installation is set at eight weeks, meaning that the lead time of the machines can never be below eight weeks. The eight-week from now period is indicated in the assembly planning by a dotted purple line.

In general, the replanning is done by moving released but not yet sold machine assembly slots forward in time to reduce the required FTEs in the week's with an FTE shortage; see Figure 2.7. An alternative is to increase the FTEs by hiring additional personnel. The goal when replanning the assemble plan is to minimise the FTE overshoot and delayed delivery.



Figure 2.7: Workload graph from week 17 onwards, left: before replanning and right: after replanning. Week 27, 28 and 29 is the summer holiday period.

2.5 Forecast accuracy

Assembly slots are released up to three months ahead during the forecast meetings, which means that machine configurations are also forecasted up to three months ahead. Therefore, it is not surprising that the reality differs from the forecast. The Operations department has determined the performance of the forecasting procedure by analysing the number of correctly forecasted modules per machine. The machine types taken into account are the machine types discussed in Section 2.1, in-between week 34, 2017 and week 13, 2021. Since week 34, 2017, the Operations department stores forecasted and eventually delivered modules in VSM's ERP system. Therefore, week 34, 2017, is chosen as the start of the study period. In week 13, 2021, the Operations department determined the performance of the forecasting procedure. Therefore week 13, 2021, is the end of the study period.

The accuracy is determined for all the considered machine types produced in this period by dividing the number of correctly predicted modules by the total number of modules considered for the machine. See Figure 2.8 for an example.

Machine A

	Module 1	Module 2	Module 3	Module 4	Module 5	Module 6	Module 7	Module 8	Module 9	Module 10
Forecasted	Yes	No	No							
Installed	Yes	Yes	Yes	No	Yes	No	No	Yes	Yes	Yes
Correctly forecasted?	Yes	Yes	Yes	No	Yes	No	No	Yes	No	No

Total number of correctly	
forecasted modules	5
Total number of modules	
considered or installed	10
Accuracy	50%

Figure 2.8: Example of accuracy determination for machine A.

10

VSM produced 295 considered machine types during the period under research. The Operations department has determined the forecasting accuracy for all of these machines. The complete histogram of the forecasting accuracy can be found in Appendix B; Figure 2.9 shows the part we think is the most relevant for visualisation purposes. After analysing the forecast accuracy, we conclude that 78% of the modules are forecasted correctly on average.

Confidential

Figure 2.9: The relevant part of the forecasting accuracy histogram, see appendix C for the total histogram.

2.6 Conclusion

This chapter answers research question 1: 'What is the current situation at VSM?'. Section 2.1 described the machine segments of VSM and determined the machining processes each product segment can perform. Besides that also the machines we need to consider in this research are determined. We analysed the assembly facility of VSM in Section 2.2 and concluded that VSM applies a fixed-position assembly layout as their machines are heavy and unwieldy while the manufacturing volumes are low. The fixed-position layout enables VSM to produce at a pace the market requires while keeping inventory levels low.

Section 2.3 determines the departments involved in the assembly process at VSM, which proved to be the Sales department that sells the machines. The Operations department that organises the machine assembly, and the Purchase department, which orders all the necessary parts. How the work procedures of these departments converge is stated in Section 2.4, which describes the assembly and purchase planning procedures of VSM.

From Section 2.4, we conclude that the lot-for-lot inventory and purchase strategy applied by VSM does not work well with the applied forecasting method and connected assembly planning procedure, as this results in a lot of replanning, FTE overshoot and long lead times. To determine the share of the forecasting procedure in this poor performance, Section 2.5 analysed the accuracy of VSM's forecasting procedure. From this analysis, we conclude that, on average, 78% of the modules are predicted correctly.



3. Literature review

In Chapter 2, we analysed the current situation at VSM regarding its assembly process. This chapter discusses the literature we reviewed to answer the second research question: 'What relevant knowledge from the academic literature can we use to determine which modules and in what quantity VSM needs to keep in stock?'. Section 3.1 starts by identifying which aspects we need to take into account when determining base stock levels. Section 3.2 covers which methods are suitable for determining the required base stock levels. Section 3.3 defines the research gap for which currently inadequate information is available, and Section 3.4 reviews the literature on the methods we could apply to close the research gap. We conclude this chapter in Section 3.4.

3.1 Aspects to consider when determining base stock levels

The decision to apply a modular product assembly strategy stems from companies' decisions on how they can best strategically align their activities with the market requirements to compete successfully. Olhager (2010) states that companies in any market need to strategically align their operations to the market's requirements to compete successfully. Companies often try to realise this by incorporating the CODP into their strategic manufacturing and supply chain operations (Olhager, 2010).

The CODP indicates the point in the product value chain where customers and products are connected. Sharman (1984) describes the CODP as the point where product configurations get frozen and the last point in the system where inventory is held. Therefore the CODP is a strategic point as the product lead times promised to the customers are based on inventory levels and production capacity available at that point in time (Olhager, 2003).

There is a broad consensus in the literature that the CODP can be differentiated into (1) make-to-stock, (2) assemble-to-order, (3) make-to-order, (4) engineer-to-order and (5) configure-to-order (Bozarth et al., 1996; Wikner and Rudberg, 2005; Donk and Doorne, 2016). Configure-to-order is a sub-strategy of assemble-to-order.

- (1) Make-to-stock (MTS): Products have a standard design, and demand is large enough to justify an inventory of finished goods to enable "off the shelf" delivery (e.g. pencils, calculators). MTS is applicable for basic consumables sold in a wide range of shops (Bozarth et al., 1996).
- (2) Assemble-to-order (ATO): Products can have a range of configurations; the configurations of the end product depends on the customer requirements. ATO is applicable for products of which the final configuration only has to be determined just before the final assembly stage, i.e. the products consist of pre-engineered and pre-assembled modules (Bozarth et al., 1996).
- (3) Make-to-order (MTO): Products consist of standard components; the configuration of the final product depends on the customer requirements. The difference with assemble-to-order is that make-to-order products tend to have unique configurations, i.e. all products are unique (Bozarth et al., 1996).
- (4) Engineer-to-order (ETO): Products are designed especially to fulfil customer requirements. Although some of the products might consist of standard components, some components are made to specific customer wishes. A typical example of ETO products is a tool and die shop (Bozarth et al., 1996).

Olhager (2010) created a clear figure based on Sharman (1984) that explains the differences between the distinguished CODP points, see Figure 3.1.

Customer order decoupling points	Engineer	Fabricate	Assemble	Deliver
Make-to-stock	Forec	ast-	>COD	$P \longrightarrow$
Assemble-to-order	drive	en	DDP Custor	mer
Make-to-order	>C(DDP	order-di	riven
Engineer-to-order	CODP			



There is also a wide consensus in the literature that upstream and downstream material flow operations differ greatly (Bozarth et al., 1996; Wikner and Rudberg, 2005; Donk and Doorne, 2006) as upstream material flow operations are forecast-driven, whereas customer orders dictate material flow operations downstream.

Section 2.1.2 indicated that the engineers of VSM designed the machines such that individual modules can independently be created, modified, replaced, or exchanged in order to enable configure to order. Therefore, we define the CODP at VSM as ATO.

3.1.1 Assemble to order

VSM applies an ATO, CODP, as the machines of VSM consist of individual modules to enable configuration to order by customers. Therefore we analyse the academic literature about modular assembly in this section.

Peas et al. (2018) distinguish two subdivisions of modular assembly, namely, modular product design (MPD) and reconfigurable modular systems (RMS). Peas et al. (2018) defined RMS as production or assembly processes that enable rapid software and hardware infrastructure changes. RMS are very suited to quickly react to market or organisational changes (Koren et al., 1999). The assembly facility of VSM consists of assembly cells explicitly designed for specific machine types, as discussed in Section 2.2. Therefore, we conclude that VSM does not have an RMS assembly process, and since VSM's management has no plans to change the assembly layout in the near future, as indicated in the research scope in Section 1.4, we do not further study the academic literature regarding RMS.

Peas et al. (2018) define MPD as products composed of several individual sub-assemblies, known as modules, which can be treated as stand-alone units yet perform as a whole. MPD is applied in many sectors, from consumer electronics to cars (Sarker and Pan, 2001). This is because modular assembly enables the mass customisation of products and services while ensuring an acceptable delivery time. Mass customisation is in great demand in the current age of globalisation (Mourtzis, 2016), which is characterised by fast-paced, complex technological developments and increasing customer demands regarding product customisation and delivery times (Mourtzis, 2016).

Pourtaleb et al. (2013) distinguish seven types of mass customisation: (1) Co-customization, (2) Custom-fabrication, (3) Assembly-by-company, (4) Assembly-by-customer, (5) On-delivery-customization, (6) Embedded customisation and (7) Standard-customization. VSM applies assembly-by-company because customers compose a machine based on the available options from the catalogue. Maalouf et al. (2020) state that an applied assembly-by-company mass customisation approach works best when pre-fabricated parts and components are available. In addition, Maalouf et al. (2020) also state that the added value for the customer lies in the assembly of the final product,



meaning that the CODP is just before the final assembly of the machine, which coincides with the previously identified ATO CODP at VSM.

The usage of modular assembly has the advantage of being able to meet customer demand better and opens up possibilities for improving the performance of the production process. Shoval et al. (2017) state that implementing a modular product assembly strategy has the following advantages for the performance of the production process: reduction of assembly costs, improved reliability, reduced assembly time and improved logistics and inventory management.

Vos (2001) argues that the successful realisation of a flexible assembly strategy for modular products requires suitable solutions to three key problems: (1) product design, (2) assembly layout design and (3) planning method. Therefore we study the literature to determine how to overcome these three key problems. Subsection 3.1.2 covers modular product design, and Subsection 3.1.3 covers assembly layout design. Subsection 3.1.4 covers the planning method.

3.1.2 Modular product design

As discussed in Section 3.1.1, Vos (2001) states that the successful realisation of a flexible assembly strategy for modular products requires suitable solutions to the three key problems: (1) product design, (2) assembly layout design and (3) planning method. This section studies the literature to determine suitable solution methods for modular product design.

All technical products have a certain degree of coupling between modules; few products consist of inseparable components. Therefore, most technical products are somewhat modular (Schilling, 2000). The success of a modular assembly approach depends on the ability to select the suitable modules to be used. Predesigned standard modules should reduce design adjustment to achieve customer satisfaction (Fujimoto and Nobeoka, 2004).

To make full use of a modular strategy, modules must be interchangeable, which means that several modules can be paired to the same base product to create products with different characteristics. We find examples of this in the automotive industry, as many components (e.g. engines, tires, radios and dashboards) can be used to create the same car model (Ulrich, 1994).

Unfortunately, due to technical limitations, complete modularity is not always achievable for complex technical products. Hölttä-Otto and Weck (2007) state that modularity is not a binary characteristic, as products can have different degrees of modularity. They prove that products subjected to many technical constraints (e.g. weight, volume, power) have a more integral architecture and that because of that translating them into a completely modular design is more challenging. Hölttä-Otto and Weck (2007) illustrate this with an example of a laptop and a desktop computer. The transformation of the laptop into a fully modular design proved to be more challenging than that of the desktop, as the design of the laptop is subject to more constraints.

Höltta-Otto and Weck (2007) also state that determining the most suitable product modules requires in-depth technical knowledge of the product, which is difficult to obtain.

3.1.3 Assembly layout design

As discussed in Section 3.1.1, Vos (2001) states that the successful realisation of a flexible assembly strategy for modular products requires suitable solutions to the three key problems: (1) product design, (2) assembly layout design and (3) planning method. This section studies the literature to determine which assembly layout designs are suitable for the assembly of modular products.

Assembly layouts have a major impact on production efficiency, as the arrangement of machines and other equipment greatly influences the required production steps and sequences. Classical classification of production layouts is (1) product line layout, (2) fixed product layout, (3) cellular layout and (4) process layout (Gayam et al., 2020).

- (1) Product line layouts: are fitting for products that need to be produced in high ongoing volumes and that require repetitive production processes. Workstations are organised in a line and perform one specific production step. The products move along the production line. Automobile and consumer electronics are typical examples of products suitable for production line assembly (Gaither and Frazier, 2004).
- (2) Fixed-position layout: If products are unwieldy, often fixed-position layouts are the best option. With fixed-position layouts, the product stays in the same location and the workers, machinery and parts are moved to the product. Ships, aeroplanes and construction projects are examples of products suitable for fixed position assembly (Gaither and Frazier, 2004).
- (3) Cellular layouts: are a combination of product line layouts and fixed position layouts. Work cells are compact production units consisting of multiple machines and operators arranged in sequential order. Work cells are composed to produce complete modules or products. Operators within a cell are usually trained to perform multiple tasks in order to make better use of the operator (Gaither and Frazier, 2004).
- (4) Process layouts: arrange the production per necessary process step (e.g., milling, cutting, punching, painting). Operators and machines that perform the same task are grouped together in a workstation. Products are transported from workstation to workstation depending on the process required for the product in question. Process layouts are suitable for firms that produce a wide variety of relatively small products in small numbers, such as customised machine parts (Gaither and Frazier, 2004).

If designed correctly, then modular product configurations have numerous assembly operations in common, which provides opportunities and challenges when designing an assembly layout (He and Kusiak, 1997).

The most discussed assembly layout in the literature regarding modular product assembly is the product line layout. Product line layouts are often designed to enable the assembly of multiple product configurations or even multiple products on the same production line in intermixed production sequences (Asadi et al., 2019). However, this is only feasible by designing the product modules so that only a limited number of assembly operations is required for the final assembly. Almost all product mixes are theoretically possible. However, the influence of the assembly sequence on the workload at the workstations and the required inventory levels is considerable, which makes designing and operating product line layouts a trade-off between minimising work overload (i.e., balancing workload) and levelling part usage (Boysen et al., 2009; Shaik et al., 2015).

Fixed position layouts are recommended for the assembly of modular products when the demand for a set of similar products is relatively low and setup times are long (Battini et al., 2007). Assembly line layouts are unsuitable for high product mix variety with low manufacturing volumes and long setup times. As the positive effects of task repeatability of product line layouts cannot be utilized due to the low manufacturing volumes and long setup times. In order to determine the most suitable assembly layout for high product mix variety with low manufacturing volumes and long setup times.



al. (2007) studied mixed-model assembly layout configuration problems with task allocation issues. These assembly layout configuration problems couple low manufacturing volumes with high product complexity (i.e., number of tasks and components) and setup times. Based on this analysis, Battini et al. (2007) conclude that fixed position layouts with a one-piece flow are often the best fit.

Cellular layouts and modular product assembly are a good match, especially with intermediate product demand and high product-mix variability. Assembly cells facilitate this by combining the positive effects of task repeatability of product line layouts and the flexibility of fixed position layouts. Cellular assembly layouts do this, when properly designed, by bringing together in one place all the required equipment, people, and processes needed for the production of a specific product family (Maalouf et al., 2020).

3.1.4 Planning methods

As discussed in Section 3.1.1, Vos (2001) states that the successful realisation of a flexible assembly strategy for modular products requires suitable solutions to the three key problems: (1) product design, (2) assembly layout design and (3) planning method. This section studies the literature to determine which planning methods are suitable for the assembly of modular products.

Modular assembly enables the mass customisation of products but requires flexibility of the planning method to deal with demand fluctuations and variability in product configuration. The planning method needs to receive real-time data from the entire supply chain about suppliers, distribution, transportation, and customer orders to achieve this flexibility. The difficulty is determining which data should be taken into account to optimise the planning method (Maalouf et al., 2020).

Most of the literature regarding mass customisation and planning focuses on determining the optimal production sequence to ensure a balanced workload and levelling of material demand for production lines. Therefore, these planning methods have dealt with so-called multi-objective optimisations problems (Rahimi-Vahed and Mirzaei, 2007; Boysen et al., 2009; Heike et al., 2001). Boysen et al. (2009) conclude after a thorough literature review that three alternative planning approaches for determining the optimal production sequence for assembly lines could be distinguished: (1) mixed-model sequencing, (2) consecutive sequencing and (3) workload-oriented level scheduling.

- Mixed-model sequencing: Aims to minimise sequence-dependent work overload by creating very detailed assembly schedules by taking into account required operational times, movement of workers and other operational characteristics affecting the assembly process (Boysen et al., 2009).
- Consecutive sequencing: Aims to create assembly schedules without requiring detailed information about the assembly process. The assembly schedules are created by applying restrictive sequencing rules. These restrictive rules ensure, for example, that no cars with consecutive special options (e.g. sunroofs and special finishes) are scheduled shortly after each other (Boysen et al., 2009).
- Workload-oriented level scheduling: Level scheduling aims to determine the assembly sequence that enables JIT delivery of parts. To achieve this, an assembly sequence that minimises the deviations between the actual and ideal assembly rates are created. Most of the literature on level scheduling focuses on the demand of materials. However, the same principle can be used for capacity utilisation (Boysen et al., 2009).

Planning methods for mixed-model assembly with low manufacturing volumes are scarcely studied in the academic literature (Heike et al., 2001). This means that the suitability of planning methods for other than production line layouts is not extensively studied. Therefore, we study the literature regarding general planning methods suitable for multiple projects instead.

Herroelen and Leus (2005) state that multi-project organisations are characterised by a high degree of complexity and uncertainty about their activities and operations. This uncertainty stems from the fact that the necessary information often only becomes gradually available, which causes blind spots in the tactical planning phase and operational uncertainty on the shop floor.

Organisations can deal with this uncertainty in two ways: proactively or reactively. The proactive approach reduces the effects of uncertainty by incorporating planning leeway to deal with unexpected events. The reactive approach mitigates the undesirable effects of uncertain events by replanning; a reactive approach is practical when undesirable events or effects are difficult to predict (Herroelen and Leus, 2005).

To support managers in selecting the most suitable planning method for a complex, multi-project environment with uncertainty, Hans et al. (2007) composed a hierarchical planning-and-control framework; see Figure 3.2. Three hierarchical decision-making levels are distinguished: (1) strategic, (2) tactical and (3) operational. And three functional planning areas: technological, resource capacity planning and material coordination.

- (1) *Strategic:* long-term is about the direction of the entire organisation, its vision, objectives and values. The strategic level is the foundation of the organisation and dictates decisions in the months to come.
- (2) *Tactical:* mid-term is about the tactics the organisation plans to use to achieve the ambitions outlined on the strategic level. The tactical level allocates the resources and creates the plans for the coming weeks.
- (3) *Operational:* short-term is about the decisions made for the organisation's day-to-day operation and the creation of schedules needed for the implementation of the plans created in the tactical phase.



Figure 3.2: Hierarchical planning-and-control frameworks by Hans et al. (2007).



3.2 Approaches for determining base stock levels

After analysing which aspects are normally considered by assembly strategies for modular products, we now analyse how base stock levels can be determined. The literature (Teunter et al.,2017; Song, 1998; Thonemann et al.,2002) states that the best method to determine base stock levels is to let them depend on the item fill rate. The fill rate is the percentage of demand that can directly be satisfied from stock-on-hand (Teunter et al.,2017). Base stock levels should not be confused with safety stock levels, as safety stock is an extra quantity of a product kept in stock to prevent an out-of-stock situation during a review period due to demand or lead time deviations.

Although the literature states that item fill rate is the best method to do this, the literature also states that there actually is minimal advice on how this can be achieved on an individual stock-keeping unit (SKU) basis in a multi-SKU environment (Teunter et al.,2017; Song, 1998; Thonemann et al.,2002). Teunter et al. (2017) state that the common practice is to set the individual item fill rates equal to the targeted fill rate of the entire inventory. Another common method is to set the individual item fill rates equal to that of a certain SKU class when a distinguishment between SKUs has been made (e.g. ABC classification).

However, these methods are inaccurate and lead to higher inventory values than necessary because not all SKUs have the same impact on the systems fulfilment performance, i.e. some SKUs are overstocked. See Figure 3.3 for a graphical representation. Overstocked inventory systems are below the efficient frontier. The challenge is to set the individual SKU levels so that the overall fulfilment performance equals the targeted fulfilment performance while keeping the inventory value as low as possible (Teunter et al., 2017; Closs et al., 2010).



Figure 3.3: System fulfilment performance VS Inventory value based on (Cardós et al., 2013).

Multiple methods to determine optimal inventory levels have been developed; we classify these models with the following three dimensions:

- 1. Product design (e.g., parts, modules)
- 2. Replenishment time (e.g., stochastic, deterministic)
- 3. Demand process (e.g., stochastic, deterministic)

Axsäter (1993) provides several algorithms for determining item fill rates under different probability distributions and ordering policies, which are unfortunately not directly applicable to practical situations. Motivated by the need to overcome this, Thonemann et al. (2002) decide to modify some

of these algorithms so that they could be used to determine optimal individual SKU inventory levels in a multi-SKU environment.

Thonemann et al. (2002) do this for an after-sales environment characterised by a Poisson demand distribution and a base stock ordering policy with order quantities equal to one (i.e., when used, a replenishment order is sent out), where all parts have the same deterministic replenishment lead time. Thonemann et al. (2002) also created a deterministic marginal analysis heuristic from practical implementation. The MA heuristic deterministically increases the inventory levels based on the demand-weighted average fill rate improvement, achieved by increasing the inventory level (i.e., 'biggest bang for the buck') until the objective is met.

Agrawal and Cohen (2001) created a multi-item inventory method to quantify the impact of component inventory policies on the delays of finished products due to component shortage. The results of this model can be used to determine optimal base stock levels for components. The model considers a make to stock policy where a sold finished product triggers a corresponding replenishment order for the components. Agrawal and Cohen (2001) state that the demand for the finished products is random and follows a known distribution, yet they do not state which distribution this is. The used replenishment lead times for the components are deterministic, and components can be back-ordered.

A simulation-based genetic algorithm for inventory level optimisation across several stages of a singleproduct supply chain is designed by Daniel and Rajendran (2005). The inventory level optimisation is based on minimising the total supply chain cost; every stage in the supply chain has its own holding and shortage costs. The genetic algorithm creates possible base stock level solutions, which are evaluated with the help of the simulation model. Daniel and Rajendran (2005) only analyse the storage of one finished product across several stages of a supply chain. The demand is assumed to be stationary by Daniel and Rajendran (2005) to reduce the complexity of the problem. The model is executed with both deterministic and stochastic replenishment times to determine the impact on the possible solutions.

Another approach is chosen by Avsar et al. (2009), who create an approximate continuous-time Markov chain queueing model for determining base stock levels. Avsar et al. (2009) only analysed one finished product consisting of two components. The assembly of the components (i.e., replenishment time) is considered and occurs at an exponential rate at single-server manufacturing facilities. The demand for finished products arrives according to a Poisson process.

Albrecht (2014) opted for a heuristic approach for determining near-optimal base stock levels to minimise long-run expected inventory and back-ordering costs for a two-stage general inventory system. Albrecht's (2014) heuristic decomposes an assembly system into multiple assembly systemsone for each end product. For each assembly system, the base stock levels are then determined. After determining the base stock levels for all the separate assembly systems, they are summed up for each component. This, unfortunately, results in overstocking of common components as the synergy effects of component commonality is not accounted for. To overcome this, Albrecht (2014) applies an echelon mapping method. The mapping method considers the demand at the final nodes and the replenishment time needed to get the components to the different echelons of the system. Based on this, the base stock levels are determined. In the heuristic, demand is fulfilled from stock consisting of finished products. The replenishment and assembly times of the components are deterministic, and the assembly capacity is deemed unlimited. Demand is stochastic, and when out of stock, demand is back-ordered.


Song and Yao (2002) have conducted an exact performance analysis and optimisation for a single product in an assemble to order system with a base stock inventory control method. The products are assembled from subassemblies that are built to stock. Song and Yao (2002) assume the time to combine these subassemblies to a finished product is neglectable, meaning that finished products are immediately available when there is inventory. Customer demand is assumed to arrive according to a Poisson distribution and the replenishment times of the components are independent and identically distributed. When analysing how base stock levels affect performance, they concluded that the problem is intractable and computationally hard; therefore, they decided to opt for upper and lower base stock level bounds that are easier computable as a surrogate solution. Song and Yao (2002) also concluded that extending to systems with multiple products turns out to be far from routine and that additional research for this is required.

3.3 Research gap

The common consensus in the studied literature is that determining optimal base stock levels is difficult and computationally tedious, especially when considering component commonality (Thonemann et al., 2002; Agrawal and Cohen, 2001; Daniel and Rajendran,2005; Avsar et al., 2009; Albrecht, 2014; Song and Yao, 2002). Due to the complexity of this problem, most developed methods are highly theoretical and simplified, making them unsuitable for practical, real-world applications. The simplifications manifest themselves in several factors. Most developed research methods only consider one machine consisting of a very limited number of components (i.e., \leq 2). Replenishment times are often deterministic, or the customer demand is, sometimes both. Assembly times are often considered at all; in the rare case that they are considered, this is done in a deterministic manner. Combinations of these factors make methods less applicable for direct real-world applications. Some of the papers state that their models can easily be extended to real-world problems; Song and Yao (2002) state otherwise, which we agree with. Therefore we think that a realistic, practical approach suitable for multi-product assemble-to-order systems with complex product configurations would be a good contribution to this field of study.

3.4 Solution approaches

To develop a practical approach suitable for determining base stock levels in multi-project assembleto-order systems with complex product configurations, we need to know what is written in the academic literature about possible solution methods and when to use or not to use them.

If it is possible, safe and cost-effective to conduct the experiments with the real physical system, it is probably desirable to do so as then there is no question about the validity of the experiments. However, this is often impossible as the experiments are too disruptive or costly (Law, 2015). When experimenting with the real physical system is impossible, one could try to find the optimal solution by complete enumeration. However, this is also often impossible for real-world instances due to their large problem sizes and solution spaces. To overcome this, mathematical algorithms can be applied (Stadtler et al., 2015).

Preferably exact mathematical optimisation methods (e.g., linear programming, simplex method, integer linear programming, branch and bound, dynamic programming or cutting plane methods) are applied, as exact methods guarantee an optimal solution (Rothlauf, 2011). However, many real-life problems are too complex (i.e., NP-hard) to be solved by exact optimisation methods; plus, the reality is often uncertain. Therefore, alternative methods, called heuristics, have been developed. Heuristics do not guarantee to find an optimal mathematical solution, but they find a solution in a reasonable amount of time, based on experience or judgment (Silver, 2014). Three different types of heuristics can be distinguished: constructive, improvement and metaheuristics.

Constructive heuristics create a complete solution from scratch by iteratively adding building blocks to a partial solution (Rader, 2010; Rothlauf, 2011). Constructive heuristics can be deterministic, probabilistic or random. Constructive heuristics do not apply backtracking (i.e., reconsider choices); therefore, constructive heuristics are 'single-pass' algorithms (Rader, 2010).

Improvement heuristics begin with a feasible solution created by a constructive heuristic and iteratively searches its neighbourhood for a better neighbouring solution (i.e., local search); see Figure 3.4 left. Neighbouring solutions are generated by altering the current solution to a certain degree. If the improvement heuristic finds a solution that it deems optimal or when a time-bound is surpassed, it terminates, meaning that an improvement heuristic results in a local optimum. A local optimum is not necessarily the global optimum; see Figure 3.4 right; this is known as the local optimality trap (Rader, 2010; Rothlauf, 2011).



Figure 3.4: left, Local neighbourhood search converging to local optimum. (Bani-Hani, 2020)

Right, local(x_1) *and global*(x_2) *optimum for search space f*(x).

Other approaches, known as metaheuristics, have been developed to overcome the local optimality trap. Metaheuristics are non-problem specific approximation optimisation strategies that 'guide' the search process to find near-optimal solutions. Metaheuristics incorporate method-specific mechanisms to avoid getting trapped in a local optimum (Blum and Roli, 2003). An often applied classification method for metaheuristics is the number of candidate solutions that are analysed simultaneously (Hu et al., 2020).

Metaheuristics that search with one solution at a time are so-called single-solution-based algorithms (SSBA). Metaheuristics that search with multiple solutions simultaneously are so-called populationbased algorithms (PBA). Metaheuristics that belong to the group of SSBA are hill-climbing, tabu search and simulated annealing. Metaheuristics that belong to the group of PBA are also known as evolutionary algorithms, such as swarm intelligence and genetic algorithms (Hu et al., 2020). SSBAs are easier to implement but have the main disadvantage that the quality of the final solution is extremely dependent on the quality of the initial solution. PBAs do not have this problem as they simultaneously search in multiple directions per iteration while exchanging information between the search agents. The exchange of information enables PBAs a better chance of escaping a local optimum and also leads to solution regions that might have better solutions (Hu et al., 2020).

3.5 Conclusion

This chapter answers research question 2: 'What relevant knowledge from the academic literature can we use to determine which modules and in what quantity VSM needs to keep in stock?'. Section 3.1 analysed which aspects are normally considered by assembly strategies for modular products. Based on Olhager (2010), we concluded that companies in any market need to strategically align their operations to the market's requirements to compete successfully. Companies often try to realise this by incorporating the CODP into their strategic manufacturing and supply chain design. We concluded that of the four distinguished CODP options (i.e., MTS, ATO, MTO, ETO), VSM applies the ATO strategy.



Vos (2001) argues that the successful realisation of a flexible ATO requires suitable solutions to three key problems: (1) product design, (2) assembly layout design and (3) planning method.

Ulrich (1994) states that to make full use of an ATO strategy, modules must be interchangeable, which means that several modules can be paired to the same base product to create products with different characteristics. Unfortunately, Hölttä-Otto and Weck (2007) concluded that complete modularity is not always achievable for complex technical products due to technical limitations. They also stated that determining the most suitable product modules requires in-depth technical knowledge of the product, which is difficult to obtain.

Gayam et al. (2020) state that assembly layouts have a major impact on production efficiency, as the arrangement of machines and other equipment greatly influences the required production steps and sequences. According to He and Kusiak (1997), this is especially the case when ATO products are designed correctly, meaning they share numerous assembly steps. There are four classical production layouts: (1) product line layout, (2) fixed product layout, (3) cellular layout and (4) process layout (Gayam et al., 2020). Battini et al. (2007) conclude that product line layouts are suitable for products with a lot of task repeatability and high production volumes, while the fixed position layouts are suitable for intermediate product demand and high product-mix variability (Maalouf et al., 2020). We concluded after studying the literature that process layouts are not suitable for the assembly of ATO products.

Maalouf et al. (2020) state that ATO enables the mass customisation of products but requires flexibility of the planning method to deal with demand fluctuations and variability in product configuration. They also state that the planning method needs to receive real-time data from the entire supply chain about suppliers, distribution, transportation, and customer orders to achieve this flexibility. The difficulty is determining which data should be taken into account to optimise the planning method. Hans et al. (2007) state that selecting the planning method can best be based on the hierarchical decision-making level (e.g., strategic, tactical, operational) and functional planning area (e.g., technological, resource capacity, material coordination).

Section 3.2 analysed the literature regarding methods for determining base stock levels. Teunter et al. (2017), Song (1998) and Thonemann et al. (2002) states that the best method to determine base stock levels is to base them on the item fill rate. However, they also state that there is minimal advice on how to do this. According to Teunter et al. (2017), the common practice is to set the individual item fill rates equal to the targeted fill rate of the entire inventory. However, this leads to overstocking as not all SKUs impact the system's performance evenly, leading to overstocking on some SKUs. The challenge is to set the individual SKU levels so that the overall performance equals the targeted performance while keeping the inventory value as low as possible (Teunter et al., 2017; Closs et al., 2010).

Several exact, heuristic, queuing, and simulation methods have been developed over the years for determining base stock levels (Thonemann et al., 2002; Agrawal and Cohen, 2001; Daniel and Rajendran,2005; Avsar et al., 2009; Albrecht, 2014; Song and Yao, 2002). The consensus of the developers of these methods is that determining optimal base stock levels is difficult and computationally tedious, especially when considering component commonality. Due to the complexity of this problem, most developed methods are highly theoretical and simplified. Therefore, we conclude in Section 3.3 that there is a lack of realistic, practical approaches in the literature suitable for multi-product assemble-to-order systems with complex product configurations.

To close the research gap, we analysed the literature in Section 3.4 about solution approaches that could be applied to achieve this. Based on Law (2015), we concluded that if possible and cost-effective,

one should always try to conduct experiments with the real system as then there is no question about validity. When conducting experiments with the real system is impossible, complete enumeration of all the possible solutions might be a suitable approach if the solution space is small enough. Unfortunately, we have to conclude that this is often not the case for real-world problems, and Stadtler et al. (2015) states that it is then best to apply a mathematical algorithm.

Rothlauf (2011) concludes that exact mathematical optimisation algorithms should preferably be applied as they guarantee to find the optimal solution. Unfortunately, he also concludes that exact methods are not applicable for complex problems that are NP-hard. Based on Silver (2014), we conclude that we can best apply a heuristic approach when a problem is NP-hard.



4. Solution design

In Chapter 3, we presented the results from the literature review. This chapter discusses the solution design we developed to answer the third research question: 'How to determine which modules need to be kept in stock and in what quantity at VSM?'. The chapter starts with the problem definition in Section 4.1. Section 4.2 covers the solution approaches we deem the most viable. The synergy between the heuristic approaches and the simulation model of VSM that we require to test them is discussed in Section 4.3. Section 4.4 analyses the heuristic approaches in more detail, and Section 4.5 does this for the simulation model. Section 4.6 concludes this chapter.

4.1 Problem definition

In Section 2.1, we discussed VSM's product segment and determined which machine types we include in this research based on the instructions of VSM's management. Therefore, the machine types that we consider per product segment are:

Beam processing	Plate processing
 Saw VB1050 	V304
 Saw VB1250 	 V310
 V600 	 V320
 Drill V630-1050 	 V325
 Drill Saw V630-1050 	FL 10 A L
 Drill V631-1050 	Flat & Angle processing
 Drill Saw V631-1050 	V550-7
 V807 	Surface treatment
V2000	VSB

The machines of VSM consist of pre-engineered modules as they are designed based upon a configure to order strategy, which allows customers to configure their machines from several pre-engineered modules incorporating different features. Appendix C shows the 288 modules for the 15 machine types considered in this research and their value in euros. Each machine consists of several fixed-base modules that are always required, some selection modules for which one of the possible choices must be selected and some optional modules that can either be added or left out; see Figure 4.1 for an example of the V807. The probability of how often customers select certain modules of the V807 is also shown in Figure 4.1.



Figure 4.1: Configure to order options V807.

Currently, VSM forecasts expected demand and machine configurations to be able to schedule ahead and order the required parts as no anonymous inventory is kept. With the base stock level strategy, this procedure belongs to the past. This is because the base stock level strategy only takes action when a machine is ordered, i.e. when VSM has received the advance payment for a machine.

After receiving the advanced payment, the machine is added to the assembly schedule by VSM's operational department. When the machine assembly begins depends on when the required assembly cell type, workforce and modules for the ordered machine configuration are available. As soon as all the required resources are available, VMS initiates the assembly of the machine. After being assembled, VSM tests the machine before preparing it for shipment.

4.1.1 Assembly and test cells

Each machine type is assigned to a specific zone; see Table 4.1. Some zones have separate assembly and test cells, while others make no distinction. Machines are moved from an assembly to a test cell in the zones with separate assembly and test cells. Otherwise, the testing occurs in the same cell as the assembly. Testing can occur in an assembly cell if a test cell is not available, but machines cannot be assembled in a test cell as the required tools are missing.

	Zone	Assembly Cells	Test Cells			
Beam processing						
VB1050	10					
VB1250	10	2	1			
V600	7	1	1			
V613	9	2	1			
Drill V630-1050	9	2	1			
Drill V631-1050	9	3	1			
Drill Saw V630-1050		Saws (VB1050) are as	ssembled in zone 10,			
Drill Saw V631-1050	9 & 10	Drills (V630/V631-1050)	are assembled in zone 9			
		Combining them	n occurs in zone 9			
V2000	8		1			
V807	4		3			
Plate processing						
V304	2		2			
V310	2		3			
V320	1	1				
V325	1	1				
Flat & Angle processi	ng					
V550-7	6	1				
Surface treatment pro	ocessing					
VSB	5	1	1			

Table 4.1: Amount of test and assembly cells per zone and machine type.

4.1.2 Assembly and test hours

The assembly and testing of the machines require a certain amount of hours and have a minimum number of assembly and testing weeks (i.e., time span); see Table 4.2. The minimum time span is caused by the fact that only a limited number of people can work simultaneously on a machine as otherwise, they would only obstruct each other. The available testing and assembly hours for 2022 are estimated to be 580 for testing and 1120 for assembly by VSM's management. When testers are not occupied with testing, they help with the assembly of the machines.



Table 4.2: The amount of testing and assembly hours required per machine type and the weeks required for assembly and testing.

Machine	Assembly hours	Assembly weeks	Testing hours	Testing weeks
Saw VB1050	31	1	78	2
Saw VB1250	31	1	78	2
V600	5	1	39	1
Drill V630-1050	85	1	43	2
Drill Saw V630-1050	179	2	141	3
Drill V631-1050	38	1	40	1
Drill Saw V631-1050	107	2	138	3
V807	80	1	60	1
V2000	2	1	20	1
V304	130	2	54	1
V310	116	2	50	1
V320	98	2	60	1
V325	150	2	160	4
V550-7	78	1	106	2
VSB	58	1	65	1

4.1.3 Base stock levels

When the assembly of a machine starts, the required modules are taken from the inventory and transported to the assembly cell. For the replenishment of the used modules, we order the necessary parts. When the parts arrive at VSM, the warehousing department temporarily stores the parts until all the parts have arrived and assembly personnel is available. After being assembled, the modules are added to the inventory. See Appendix C for the required assembly time per module type.

When one of the required resources is not available, the assembly of a machine cannot start. Therefore all these resources must be in tune with each other to keep costs low and efficiency high. The number of assembly and test cells is fixed as well as the available workforce. This leaves the base stock levels as the sole parameter that we can alter in this study to reduce machine lead times.

With that, the problem comes down to determining the base stock levels needed to ensure that at least 95% of the machines of future demand scenarios can be completed within 10 weeks, given VSM's current assembly capacity (i.e., cells, personnel). When determining these base stock levels, we must ensure not to overstock as we want to keep the inventory value as low as possible. Stocking one module for all 288 module types already results in an inventory value of 5 million euros. We have to consider several constraints when determining the optimal base stock levels.

Constraints

- All demand must be fulfilled;
- The assembly of a machine cannot start before all the required modules are present;
- The assembly of a module cannot start before all the parts have arrived;
- The required amount of assembly and test hours in each week cannot exceed the capacity;
- The required number of assembly and test cells per machine type per week cannot exceed capacity;
- The fraction of a machine being assembled or tested in a week cannot exceed the maximum fraction that can be completed within a week for that machine type;
- All base stock levels must be non-negative.

To determine the optimal base stock levels, we also need to make several assumptions.

Assumptions

- The inventory levels are topped up to the defined base stock levels at the beginning of each year, i.e. at the beginning of each simulation run;
- There is an infinite inventory capacity;
- No assembly cells are required for the assembly of the modules (i.e., infinite capacity);
- Modules have no minimal assembly duration (i.e., modules can be assembled within a week);

- Mechanics can assemble all the machine types and modules;
- Testers can test all the machine types;
- When testers are not occupied with testing, they assist with the assembly of the machines;
- Assembly and test durations are deterministic;
- The assembly and test workforce per week (i.e., assembly and test hours) are constant during the year.

4.2 Solution approach

VSM sells on average 160 machines per year, which consist on average of 20 modules, resulting in a large problem size of about 3200 modules that we need to schedule for every one-year demand scenario we analyse. Besides a large problem size, we also have a huge solution space because the only restriction VSM enforces on possible base stock levels is that they need to be non-negative.

Experimenting with the real assembly system is out of the question, as VSM's management considers conducting experiments with the real assembly system too expensive, disruptive and time-consuming. Therefore we need to find another method to determine the required inventory levels. Finding the optimal solution by enumeration is also not possible due to the large problem size and solution space. Therefore we need to apply mathematical algorithms to overcome this (Stadtler et al., 2015).

Preferably we would apply an exact mathematical optimisation method; however, computing robust base stock levels is proven to be an NP-hard problem by Bienstock and Özbay (2008). The reason that computing robust base stock levels for the situation under study at VSM is NP-hard is due to demand uncertainty, which robust base stock levels have to be able to deal with for many possible future scenarios. Calculating the performance of base stock levels for many future scenarios with, for example, dynamic programming is extremely time-intensive. Therefore, applying an exact mathematical method for determining robust base stock levels would only work for 'toy problems' (Rothlauf, 2011). As the problem at VSM is not a 'toy problem,' we need to apply heuristic procedures instead.

The first heuristic approach we suggest is a population-based algorithm (PBA) and the second approach is a local search procedure. VSM's management is interested in which of the two approaches results in the most cost-effective set of base stock levels capable of fulfilling the objective.

Approach 1: as we expect the solution space to have many local optima, we choose a PBA because of its ability to search through solution spaces with multiple local optima. We deem the genetic algorithm (GA) the most suitable PBA for the problem instance at VSM. An important reason for this choice is the GA's flexibility to deal with various objective functions while requiring a minimum of fine mathematical properties (Gen and Cheng, 2000). Other appealing features of GAs are their ability to deal with real-life size problems and their ability to use historical data to guide the search to the best performing region within the solution space (Daniel and Rajendran, 2005).

GAs are a group of computational optimisation models based on natural evaluation processes that create new solutions based on the previous generation of solutions and their performance. GA encode a potential solution to a problem in a chromosome-like data structure and apply mating and mutation (i.e., recombination) operators to create new potential solutions that preserve the critical information of good scoring individuals of the previous generation (Whitley, 1994).



Approach 2: Unfortunately, GAs, as used in approach 1, are not well suited for fine-tuning solutions (i.e., local search), which are very close to optimal ones (Martinez and Lozano, 2007). Therefore, we choose a GA and local search combination for the second approach; we refer to this approach as the local search approach. The GA guides the search to the best performing region within the solution space, after which the local search methodological searches through this region for the best solution. We choose for local search as it has proven to be very successful in determining near-optimal and sometimes even optimal solutions for difficult real-life problems with enormous solution spaces (Aarts and Lenstra, 2003; Dumitrescu and Stützle, 2003).

The GA creates the initial solution, and the local search algorithm searches through the solutions in its neighbourhood by applying small local changes to the solutions until a found solution is deemed optimal; i.e. the percentage of machines completed within 10 weeks is at least 95%.

To test the generated sets of base stock levels from both approaches, we use a simulation model as determining the performance of the generated sets with VSM's real assembly system is not practical due to disruptions, costs, and time durations. Therefore, we apply a simulation model which simulates VSM's purchase, inventory management and assembly process. The simulation model is as realistic as possible in order to function as a digital twin of the real system. The performance of the sets of base stock levels given as input to the simulation model is measured in the average percentage of machines completed within 10 weeks for multiple future demand scenarios. The synergy between the approaches and the simulation model is discussed in Section 4.3.

4.3 Synergy between heuristic approaches and the simulation model

In this section, we discuss how the two heuristic approaches and the simulation model of VSM's purchase, inventory management, and assembly process work together (i.e., synergy) to generate sets of base stock levels that with 95% certainty can guarantee that at least 95% of the machines of future demand scenarios can be completed within 10 weeks.

The two heuristic approaches generate sets of base stock levels and send them to the simulation model to determine their performance. The simulation model determines the performance of a set of base stock levels for a one-year stochastically generated demand scenario by generating an assembly schedule based on the demand scenario and the inputted set of base stock levels. Based on the assembly schedule, the lead times of the machines are determined, followed by the percentage of machines completed within 10 weeks.

We discuss the synergy between the GA heuristic and the simulation model in more detail in Section 4.3.1. The synergy between the local search heuristic and the simulation model we discuss in further detail in Section 4.3.2. We analyze the heuristics themselves in detail in Section 4.4 and the simulation model we evaluate in more detail in Section 4.5.

4.3.1 Synergy GA and simulation model

This section explains how the GA approach and the simulation model work together to generate sets of base stock levels that can meet with 95% certainty the objective of completing at least 95% of the machines of future demand scenarios within 10 weeks; see Figure 4.2 for the flow chart of this collaboration.



Figure 4.2: Flowchart of the collaboration between the GA approach and the simulation model.

This approach starts with the GA generating a set of base stock levels for a one-year demand scenario that allows at least 95% of the machines of that demand scenario to be completed within 10 weeks. How the GA does this is discussed in Section 4.4. The just determined set of base stock levels is tailored to the generated one-year demand scenario. Therefore, we need to examine how it performs for other possible future demand scenarios; for this, we use the simulation model. The simulation model creates stochastic demand scenarios (i.e., possible future demand scenarios) to determine the performance of the set of base stock levels. How the simulation model creates these stochastic demand scenarios is explained in Section 4.5.

The simulation model measures the performance of the set of base stock levels in the percentage of machines completed within 10 weeks for a given demand scenario. The simulation model performs simulation runs (i.e., schedules scenarios) for at least 10 runs or until the relative error of the performance of the set of base stock levels is less than 5%. Next, we take the average performance of all the executed simulation runs for the tested set of base stock levels. We store the set of base stock levels as a solution if the average percentage of machines completed within 10 weeks for the scheduled demand scenarios is at least 95%. However, it should be noted that this never occurs for the initial set of base stock levels.

If the average performance is below 95%, we return the set of base stock levels to the GA, which boosts (i.e., increases) certain base stock levels; how the GA does this is explained in Section 4.4. After boosting the base stock levels, we return the set of base stock levels to the simulation model to reexamine the performance for multiple stochastically created demand scenarios. The simulation model once again performs simulation runs until the relative error of the performance of the set of base stock levels is below 5%. We repeat this process of boosting the set of base stock levels until the set can meet the objective of completing at least 95% of the machines of future demand scenarios within 10 weeks.

After storing a set of base stock levels capable of meeting the objective in the solution database, we check whether we have gathered enough samples (i.e., sets) to be able to guarantee that we found a good solution. We need to gather multiple solutions as we do not know the optimal solution or any near-optimal solution (i.e., relaxed solutions). How many samples we need to gather depends on the number of decision variables we take into account (Thengvall, 2019); the larger the number of decision variables, the larger the number of samples; see Figure 4.3.



# Decision Variables	Minimum number of simulation trials
Less than 10	100
Between 10 and 25	500
Between 25 and 100	2,500
Greater than 100	5,000

Figure 4.3: The required minimum number of simulation trials per amount of decision variables (Thengvall, 2019).

Because we need to determine the required base stock levels for 288 different module types, we need to gather at least 5000 samples (i.e., simulation trials). However, as we have a double objective of meeting the required performance while also minimising the inventory value, we need to determine the optimal tradeoff between them; an often applied practical rule of thumb is to multiply the required samples (i.e., simulation trials) from Figure 4.3 by five (Thengvall, 2019). This means that we need to gather at least 25,000 sets of base stock levels in order to be able to guarantee that we found a good solution.

4.3.2 Synergy local search and simulation model

This section explains how the local search approach and the simulation model work together to generate sets of base stock levels that meet with 95% certainty the objective of completing at least 95% of the machines of future demand scenarios within 10 weeks; see Figure 4.4 for the flow chart of this collaboration.



Figure 4.4: Flowchart of the collaboration between the local search approach and the simulation model.

The first steps of the local search approach are the same as for the GA since it uses the GA to create an initial set of base stock levels for a one-year demand scenario that allows at least 95% of the machines of that demand scenario to be completed within 10 weeks. The performance of this set of base stock levels is also determined with the help of the simulation model, which determines the average percentage of machines completed within 10 weeks for multiple future demand scenarios. Until now, there are no differences between the two approaches.

However, contrary to the GA, If the average performance is below 95%, we do not return the solution to the GA, but instead, we apply a local search on the initial set of base stock levels in order to boost (i.e., improve) it; how the local search operates we explain in Section 4.4. After boosting the base stock levels, we return the set of base stock levels to the simulation model to determine the performance for multiple future demand scenarios. This process is repeated until the created set of base stock levels can meet the objective.

After storing a set of base stock levels capable of meeting the objective, we check whether we have gathered enough samples (i.e., sets). The number of samples to be gathered is the same as for the GA since the number of decision variables and the objective function are the same. After gathering enough samples, the model terminates and returns the most cost-effective set of base stock levels found.

4.4 Heuristic approaches

This section analyses the two heuristic approaches we apply to generate sets of base stock levels that are with 95% certainty capable of meeting the objective of completing at least 95% of the machine of future demand scenarios within 10 weeks —the GA approach we discuss in Section 4.4.1 and the local search in Section 4.4.2.

Section 4.3 explained that we apply the simulation model of VSM's purchase, inventory management, and assembly process to determine what machine lead times can be achieved given a certain demand scenario and set of base stock levels. However, we also use this simulation model in the GA and local search approach to determine the performance of unfinished (i.e., ongoing) sets of base stock levels to provide insight into their performance during the construction process.

4.4.1 Genetic algorithm (GA) approach

The GA's objective is to create a set of base stock levels for a one-year demand scenario that enables at least 95% of the machines of that demand scenario to be completed within 10 weeks, as explained in Section 4.3.1.

The GA does this by creating an initial population consisting of several individuals as is in nature. In GAs, the initial populations are often randomly generated (Whitley, 1994). Therefore, we also use a random approach when creating an initial population. The individuals in the initial generation are all sets of base stock levels, which means that all the individuals contain the base stock levels of the 288 distinguished module types. The generation of the initial population occurs per one individual at a time. For each individual, we randomly increment several module types by one. As the incremented module types are randomly selected, there exists a probability that a module type is selected multiple times; however, this is no problem as the initial base stock levels are far below the eventually required ones. The probability of this occurring depends on the number of incremented modules relative to the 288 module types considered in this research.

Next, the performance of all the individuals in the created initial population is determined by conducting simulation runs. The performance of the individuals is determined by simulating them all with the same input one-year demand scenario in order to compare their performance with each other. After determining their performance, we check if one of the individuals (i.e., set of base stock levels) can fulfil the objective; if so, we flag the individual as a possible solution, and the GA stops. However, it should be noted that this never occurs for the initial individuals.

When none of the individuals can fulfil the objective, we transform their performance into mating probabilities for the creation of the next generation. One of the benefits of GAs is that almost everything can be used as a function for the mating probability (Sivanandam and Deepa, 2008). We combine the inventory value and performance of the individuals in the mating probability function; see Figure 4.5.



 $P_i = performance individual$ $P_I = performance population$ $C_i = cost individual$ $C_I = cost population$

> 1 if above 0, then multiply with P_i to determine the mating probability of individual *i*. if below 0, then the mating probability is 0, as individual *i* performs below average.

Figure 4.5: Applied mating probability function.

The determined mating probabilities of the individuals indicate the chance that they are selected to create offspring (i.e., new individual). The individual's chromosome-like data structures contain the base stock levels of the 288 module types we take into account. These data structures can be cut into machine-specific sections see Figure 4.6. Because we can cut the data structures into machine-specific sections, we consider it sensible to build up new individuals per machine type to increase the GAs improvement speed. Therefore, we apply a k-point crossover approach, which cuts the string of module types of the individuals into machine-specific sections. Because we want to build up new individuals per machine type, we determine the mating probabilities per machine-specific section of an individual rather than the overall mating probability.

[Module number, Inventory level]	[1,3], [2,4], [3,5], [4,2], [5,3], [6,4],	[7,1], [8,6], [9,3], [10,2],	[], [288,4]
	Machine type 1	Machine type 2	Machine type
		C.1. 1. 11. 1. 1. 1.	

Figure 4.6: Schematic representation of the chromosome-like data structures of the individuals in a generation.

After cutting the chromosome-like data structures for each individual in a generation into machinespecific sections and determining their mating probability, we can create new individuals; see Figure 4.7. Since we build up new individuals per machine type, they can have 15 different 'parents', as we consider 15 different machine types.



Figure 4.7: k-point crossover, k can be any non-negative integer.

In nature, genetic mutations can occur, introducing diversity into the population. Diversification is wanted to reduce genetic drift. Genetic drift means that the population loses variety and character traits (Eiben and Smith, 2003). Mutations are generally simulated using mutation probabilities, which determine for each chromosome (i.e., module type) in an individual's DNA string whether or not a mutation occurs. Mutating probabilities are often less than 1%; however, in cases where a mass mutation is wanted, they can be much higher, even up to 100% (Eiben and Smith, 2003).

In order to introduce diversity into the individuals of a generation, we therefore also apply mutations to the base stock levels (i.e., increment by one) to enable the GA to explore the solution space more quickly. To improve the diversification speed of the GA, we do not apply a low mutating probability but rather a fixed number of mutations per individual. The base stock levels that are increment by one are randomly selected based on their module shortage probability relative to the overall shortage, which we determine when executing simulation runs to determine the mating probability.

After creating a new generation based on the previous generation, we determine the performance of all the individuals in the new generation by conducting simulation runs with the input demand scenario. After which we check if one of the individuals can fulfil the objective. If not, we determine the mating probabilities again and repeat the process. The idea is that the individuals in each generation become more and more capable of fulfilling the objective.

After creating a possible set of base stock levels that fulfils the objective, the GA stops, and the performance of the created set of base stock levels is determined for possible future demand scenarios with the simulation model as discussed in Section 4.3.1. The simulation model performs simulation runs, i.e. schedules future demand scenarios with the set of base stock levels generated by the GA until the relative error of the performance of the scheduled scenarios is below 5%. After that, we determine the average performance of all the executed simulation runs. We consider the average performance as the ability of the set of base stocks to deal with future demand scenarios.

If the average performance is less than 95%, we consider the ability of the set of base stocks levels to cope with future demand scenarios insufficient. When insufficient, we send the set back to the GA to be boosted. The boosting process starts with the creation of a generation consisting of several individuals by cloning the returned set of base stock levels. When cloning, we apply mass mutation by incrementing a fixed number of base stock levels by one. The base stock levels that are increment by one are randomly selected based on their module shortage probability relative to the overall shortage.

After cloning, the set of base stock levels is sent back to the GA, a new demand scenario is generated for the GA, and the performance of all the clones is determined for that scenario with the help of the simulation method. Next, we determine the mating probabilities of the clones and create a new generation. From this point on, the creation of new individuals and generations occurs in the same manner as discussed before. This process is repeated until one of the individuals can ensure that at least 95% of the machines of the new demand scenario can be completed within 10 weeks. After which we sent the set of base stock levels to the simulation model to determine its performance for multiple future demand scenarios. When its performance is insufficient, it is returned once more, and otherwise, it is stored in the solution dataset.

This process of returning sets of base stock levels with insufficient performance for future demand scenarios to the GA works because the GA creates new individuals out of machine-specific sections of previous individuals and because the base stock levels of all 15 machine-specific sections become stronger each generation through mutations. This process repeats itself until all machine-specific sections are robust enough to ensure that the created sets of base stock levels can complete at least 95% of the machines of future demand scenarios within 10 weeks.

Figure 4.8, left, shows the flow chart of the GA we apply to create the initial set of base stock levels and Figure 4.8, right, shows the flow chart of the GA we apply when improving on an insufficient set of base stock levels.





Figure 4.8, left: flow chart of the GA we designed to create an initial set of base stock levels. Figure 4.8, right: flow chart of the GA we designed to improve (i.e., boost) on an insufficient set of base stock levels.

4.4.2 Local search approach

GAs are very suited for guiding the search to the best performing region within the solution space but less for fine-tuning a solution close to the optimum (Martinez and Lozano, 2007); therefore, we decided to apply a local search approach in combination with the GA discussed in Section 4.4.1.

This approach starts with the GA generating a set of base stock levels for a one-year demand scenario that allows at least 95% of the machines of that demand scenario to be completed within 10 weeks. After this, the simulation model examines the performance of this set of base stock levels for future demand scenarios in the same manner as earlier discussed. After analysing the performance, it is determined if the set of base stock levels can ensure that at least 95% of the machines of future demand scenarios can be completed within 10 weeks. When this is the case, the solution is stored, and the algorithm stops. However, it should be noted that this never occurs for the initial set of base stock levels. Until now, there are no differences between the previously described GA and the local search approach.

The difference between the two approaches is that we do not send the set of base stock levels back to the GA to be boosted when the average percentage of machines completed within 10 weeks for future demand scenarios observed by the simulation model is lower than 95%; instead, we conduct a local search on the underperforming set of base stock levels.

The local search consists of several steps. The algorithm starts with determining how often the different module types are out of stock. This is determined by conducting multiple simulation runs with the set of base stock levels. After that, these determined out-of-stock occurrences per module type are transformed into module shortage probabilities relative to the overall shortage. However, now we only take into account the modules that are out of stock the most, and while this is important, we also need to ensure to minimise the overall inventory value.

The approach we selected to also take into account the overall inventory values starts with randomly selecting several module types based on the determined module shortage probabilities. Next, we transform the value of the sampled modules into probabilities relative to the overall value of the sampled modules. We then randomly select one or more modules to be boosted based on these probabilities. When selecting more than one module to be incremented, the probability of overshooting exists; however, it also reduces the required running time of the algorithm. By taking random samples, we ensure the stochastic character of the local search.

After boosting the set of base stock levels, we determine its performance for future demand scenarios with the help of the simulation model. If the upgraded set of base stock levels does not perform better, we do not keep the boosted set and redo the process with the unboosted set. In case the upgraded set of base stock levels performs better, we check if it can meet the overall objective of completing at least 95% of the machines of future demand scenarios within 10 weeks. If yes, we store the set of base stock levels into the solution database, and if not, we repeat the process of boosting the set of base stock levels with the local search algorithm until it does. See Figure 4.9 for the flowchart of the applied local search algorithm, which begins after receiving a not good enough set of base stock levels.





Figure 4.9: Flow chart of the Local search algorithm we applied to improve insufficient sets of base stock levels.

4.5 Simulation model

In this section, we analyse the simulation model of VSM's purchase, inventory management and assembly process we designed to determine what machine lead times can be achieved with a set of base stock levels for either a single one-year demand scenario or multiple future demand scenarios. The aspects we analyse in more detail are the demand scenario generator, the scheduling algorithm, and the model's validity.

4.5.1 Scenario generation

VSM's management has indicated that they would like to consider the stochastic effects of demand and machine configurations on the number of modules that need to be kept in stock. This is because they regularly experience these stochastic effects, which sometimes have far-reaching consequences for the assembly process and thus the lead times of the machines. Besides these stochastic effects, VSM's management also wants to consider the seasonality of demand.

The easiest way to take the demand and seasonal patterns of the different machine types into account would be by using historical data that contains these stochastic factors and seasonal patterns. However, using historical data has a drawback as it only describes what happened in history, not what could have happened. Besides this, the required historic data needed for this is not stored in an easily accessible and trustworthy way at VSM, making the creation of accurate historic datasets very time-consuming. Another drawback of using historic demand is that it does not enable the generation of many scenarios as only limited information is available. Therefore we decided not to use historic demand but to fit a theoretical distribution function to the observed data to enable the simulation model to generate an unlimited number of possible scenarios. Another advantage of using a theoretical distribution function is that now scenarios can be created with values outside the range of the historic data.

To determine the theoretical distributions that best fit the demand patterns of the different machine types, we had to identify their demand patterns. We did this by counting the demand per machine

type per month; we analysed the sales data from 2015 to 2020. Now that the demand patterns for every machine type per month are known, we determined the fit of different probability distributions to these demand patterns. The Poisson distribution had a good fit with the demand patterns of all machine types; see Appendix D. Law (2015) indicates that the Poisson distribution gives the number of events that occur in an interval of time when the events are occurring at a constant rate. An event, in this case, is the selling of a machine. Applying the Poisson distribution smoothens out the observed event frequency, which is useful as we only have a limited number of observations. We can sample the expected demand per month from the Poisson distribution by applying inverse transform sampling. Per machine type, we take a sample for each month of the year, which we add up to determine the annual demand.

Now that we can generate annual demand, we need to decide in which week this demand occurs, which depends on seasonal influences. To incorporate seasonality in the generated scenarios, we determine the number of machines ordered in every week {1..52} of the year; for this, we used the same sales data from 2015 to 2020. Based on this, we determine the probability that a machine is ordered in a certain week. From these probabilities, we sample per machine to assign it a specific order week. The determined seasonal factors are stated in Appendix D.

The last thing to determine in order to be able to create demand scenarios is how to take into account machine configurations. We decided to determine how often customers choose the different modules types per machine type. For this purpose, we analysed the same sales data from 2015 to 2020 that we used for demand and seasonality. However, as this sales data only gave a partial picture, we also consulted the sales agents of VSM for confirmation. When creating stochastic machine configurations, the simulation starts by assigning all the machines the fixed modules of their corresponding machine type. After that, we randomly assign them the optional modules based on the probability of occurrence.

The flow chart on how the simulation model generates a stochastic demand scenario is stated in Figure 4.10.



Figure 4.10: Process for generating stochastic demand scenarios.

4.5.2 Scheduling algorithm

We need to create assembly schedules based on the set of base stock levels being examined and the available assembly resources of VSM (i.e., personnel, assembly and test cells) to determine what machine lead times can be achieved with the set of base stock levels. We need to create assembly schedules to do this because there are multiple interactions between the machines to be scheduled and the available resources, as they have to be shared.

Currently, VSM's assembly schedules are made by hand by a planner, who applies some fuzzy scheduling rules and a lot of personal insight, which we, unfortunately, cannot mimic due to the number of decisions made. Therefore, as an alternative, we studied the literature to find a scheduling method (i.e., heuristic) that would generate similar schedules for a demand scenario as the planner



would. Together with VSM's management, we decided that a simple priority rule-based scheduling heuristic would best represent the actual scheduling process.

We do not determine the optimal scheduling method because that is not the purpose of this study and is therefore beyond its scope. The objective of the simulation model is to imitate VSM's scheduling processes as truthfully as possible to determine the machine lead times that VSM can achieve when a base stock level strategy is applied without further changes to the processes of VSM.

The priority rule-based scheduling heuristic divides the scheduling process into separate scheduling stages. Normally, a scheduling stage consists of a (1) remaining set, (2)decision set, (3) active set and (4) completed set (De Boer, 1998). For the scheduling situation at VSM, we add a (5) replenishment set; see Figure 4.11. The scheduling heuristic runs until all machines in a demand scenario are completed; in each scheduling stage, the characteristics of the model (i.e., sets) change as the model is dynamic.



Figure 4.11: Schematic view of priority rule-based heuristic with explanations of the different sets.

At the start of the scheduling run, i.e. first scheduling stage, the scheduling heuristic loads all the machines of the input scenario into the remaining set. The scheduling heuristic checks whether the order week (i.e., release week) of the machines in the remaining set is reached in the other scheduling stages. The scheduling heuristic transfers the machines to the decision set when this happens.

Before the scheduling heuristic starts determining which machines form the decision set to add to the active set (i.e., assembly schedule), it determines whether there are machines in the active set that are already assembled, tested and prepared for transport at that scheduling stage. Because these machines still occupy a test or assembly cell, the scheduling heuristic transfers these machines to the completed set, freeing up the cells. Then the scheduling heuristic checks the active set for fully completed modules and removes them from the active set after incrementing the respective inventory level.

After determining the inventory levels and available assembly and test cells at that scheduling stage, the scheduling heuristic checks if it has to move assembled machines from an assembly to a test cell or not. Next, the scheduling heuristic reduces the available assembly and test hours at that stage with the number of hours required for the machines and modules currently in the active set. After determining the leftover capacity, the scheduling heuristic checks for the machine with the earliest due date in the decision set if it can be added to the assembly schedule (i.e., active set). Then the scheduling heuristic checks it for the machine with the earliest due date after that, until all machines in the decision set are checked. When the scheduling heuristic transfers a machine to the active set, the required assembly cell and assembly hours are reserved. Besides that, the inventory levels of the module types needed for the transferred machine are reduced by the required quantity and added to the replenishment set.

The modules in the replenishment set get a stochastic replenishment time assigned, indicating in which scheduling stage the parts for this module arrive at VSM. A triangular probability distribution, see Figure 4.12, generates these replenishment times based on the minimum, maximum and median replenishment times for that module type; see Appendix C.



Figure 4.12: Triangular probability distribution and formulas (Law, 2015).

After adding machines to the active set, the scheduling heuristic checks the replenishment set and determines the fill rate of the modules for which the parts have arrived at that scheduling stage. Based on these fill rates and the available resources left, modules are selected for assembly and transferred to the active set. The required resources for the transferred modules are reserved.

When all machines of the input scenario are assembled, the heuristic stops and the achieved machine lead times are determined based on their order and finish data in the assembly schedule. The created assembly schedules can be visualised by a Gannt chart, see Figure 4.13, for validation or gaining more insight into the performance of a set of base stock levels. The flow chart of the scheduling process just described is shown in Figure 4.14.



Figure 4.13: Gannt chart of a one-year assembly schedule for VSM generated by the priority rule-based scheduling heuristic.





Figure 4.14: The priority rule-based scheduling heuristic we designed to create assembly schedules representative of VSM's scheduling process

4.5.3 Validation

To ensure that the sets of base stock levels generated by both approaches are really capable of ensuring for the real-world situation at VSM that with 95% of certainty at least 95% of the machines of future demand scenarios can be completed within 10 weeks, we need to guarantee the validity of the simulation model.

The best method to determine the validity is to compare the simulation models output to results observed in reality (Law, 2015). Therefore, the best method to determine the simulation models validity would be to execute it with a real-world scenario and compare the machine lead times achieved by the simulation model to the really achieved machine lead times. Unfortunately, the data needed to create a real-world scenario to test the validity of the simulation model is not stored in an easily accessible and trustworthy way at VSM. Therefore, we analyse and validate as many separate aspects of the simulation model as possible as an alternative. For the analysis and validation of the separate sections, we use the experience of VSM's management. However, it should be noted that because we cannot validate the simulation model as a whole, the validity of the simulation model is the major weak spot of the analysed approaches.

We created one thousand scenarios to validate the generated demand scenarios and determined the average number of machines ordered per machine type. These averages we compared to VSM's sales prediction for 2022. According to VSM's management, the average demand for the Drill V630-1050, Drill/Saw V630-1050, V807 and V304 were too far off to be acceptable. After analysing potential causes, we discovered that this is caused by changing demand patterns. Changing demand patterns mean that we cannot use historic sales data to predict future demand as they deviate from each other.

The changing demand patterns for the Drill V630-1050, Drill/Saw V630-1050 and V304 are caused by the fact that these machine types are phased out in the coming years by VSM. This means that the sales agents are no longer recommending these machine types to potential customers as frequently as before, which is why sales are down compared to historical sales figures. On the other hand, the demand pattern observed in the historical data for the V807 is too low, according to VSM management, because the demand for this machine is increasing rapidly, making the scenario generator too conservative in its amount of predicted sales. To address these changes in demand patterns, we have manually changed the lambda of the Poisson distribution for these machines based on the remarks of VSM management.

To validate the seasonal patterns within the demand scenarios created by the scenario generator, we analysed several created scenarios with VSM's management. We observed a gradually declining demand pattern during the year. VSM's management recognises this pattern as companies invest, i.e. reserve money on the balance sheet, to reduce taxes at the end of the year. However, these machines do not arrive as orders for the assembly department before the beginning of the following year, as the Sales department must first configure the machines based on the customer's requirements.

The machine configurations are the last aspect of the created demand scenarios to be validated. We validated this by generating several scenarios from which we have taken some machines per machine type. VSM management examined the configurations of these sampled machines to determine if they were correct. Since they did not find any irregularities, we conclude that the generated demand scenarios reflect reality.

After validating the demand scenario generator, we asked VSM's management to validate the scheduling heuristic. We did this by assessing some scenarios and the assembly schedules created for these scenarios with VSM's management to determine if they would make similar assembly schedules. After the analysis, VSM's management expressed confidence in the scheduling method. However, whether the applied scheduling heuristic accurately represents VSM's actual scheduling process is debatable.

4.6 Conclusion

This chapter answers research question 3: 'How to determine which modules need to be kept in stock and in what quantity at VSM?'. Section 4.1 defined the problem, and we concluded the problem comes down to determining the base stock levels needed to ensure with 95% certainty that at least 95% of the machines of future scenarios can be completed within 10 weeks, given VSM's current assembly capacity. We also identified the assumptions and constraints to consider when determining the base stock levels.

Section 4.2 discusses the solutions approach, and we decided that conducting experiments with VSM's real assembly system is not practical. We also decided that enumeration of all possible solutions is not practical due to the enormous solution space. Because Bienstock and Özbay (2008) proved that determining robust computing base stock levels is NP-hard, we, unfortunately, had to conclude that



exact optimisation methods are also not applicable. Therefore we decided that applying heuristic approaches is the most promising. We came up with two heuristic approaches that we deem the most viable.

The first approach is a GA; we concluded that GAs might be suitable for determining robust sets of base stock levels because of their ability to deal with real-life size problems and their ability to use historical data to guide the search to the best performing region within the solution space (Daniel and Rajendran, 2005). The second approach is a combination of a GA and local search algorithm; we refer to this approach as the local search approach. We opted for this combination as GAs are not well suited for fine-tuning solutions (i.e., local search), which are very close to optimal ones (Martinez and Lozano, 2007). The GA guides the search to the best performing region within the solution space, after which the local search takes over to analyse this region. We selected local search as it has proven to be very successful in determining near-optimal and sometimes even optimal solutions for difficult real-life problems with enormous solution spaces (Aarts and Lenstra, 2003; Dumitrescu and Stützle, 2003).

We concluded in Section 4.3 that to determine the performance (i.e., percentage of machines completed within 10 weeks for future scenarios) of the sets of base stock levels; we need to apply a simulation model of VSM's purchase, inventory management, and assembly process as conducting experiments with the real assembly process is not possible.

The two heuristic approaches we designed for determining robust base stock levels capable of meeting the objective are described in Section 4.4. In this section, we concluded that the GA could best consider both the performance and the inventory value of an individual (i.e., set of base stock levels) when determining the mating probability of that individual. We also decided to apply k-point cross over and mass mutation when creating new individuals. For the local search approach, we decided that the best approach was to randomly sample several modules based on their out of stock occurrences and then randomly select some of them based on their value relative to the overall value of the sampled modules. After which, the base stock levels of the selected module types are incremented by one.

We discussed the design of the simulation model in Section 4.5. When designing the simulation model, we decided that we needed to create realistic demand scenarios in order to test the performance of sets of base stock levels. We concluded that we have to consider historic demand, seasonality, and ordered machine configurations to make demand scenarios realistic. Together with VSM's management, we also concluded that a simple priority rule-based scheduling heuristic would best represent the actual scheduling process as the current manual scheduling process is impossible to simulate properly.

5. Analysis

In Chapter 4, we presented two approaches that we expect to be suitable for determining robust base stock levels that, with a certain degree of certainty, can meet a lead time objective for a percentage of the orders while minimising the inventory value. To answer the fourth research question: 'What is the best approach for determining base stock levels at VSM?' we analyse these two approaches. As both approaches are heuristics and no exact approaches and because enumeration is also not possible due to the sheer size of the solutions space, we, unfortunately, cannot determine if a found solution is the optimal solution. Therefore we use the two heuristic approaches as a benchmark for each other.

We analyse the procedure applied by both approaches for the generation of sets of base stock levels in Section 5.1. The data sample sizes required for further analysis of both approaches are determined in Section 5.2. The lead time performance and the inventory values of the sets of base stock levels generated by the potential approaches are analysed in Section 5.3, and Section 5.4 discusses if we found a suitable solution for the research gap and thus for VSM. Section 5.5 concludes this chapter.

5.1 Procedure for generating sets of base stock levels

From the description of both approaches in Section 4.3, it can be deduced that both work on the basis of starting from a set of base stock levels that are well below the required minimum base stock levels and then increasing these base stock levels until the performance that can be achieved with the base stock levels meets the objective. We will, from now on, refer to this as the incrementing approach. How the base stock levels are incremented is different for both approaches, yet both approaches do this randomly based on the performance of the set of base stock levels for several stochastically generated future demand scenarios. As a result, the probability of selecting a certain base stock level (i.e., module type) for incrementing is different in each stage of the improvement process. Due to this and the fact that we take 288 different module types into account, the approaches in principle always create unique sets of base stock levels; we did not yet observe two identical sets of base stock levels.

Another possible approach would have been setting all the base stock levels to the same predetermined maximum level and determining if they could be reduced while still meeting the objective of ensuring that with 95% certainty, 95% of the machines of future demand scenarios can be completed within 10 weeks. We will, from now on, refer to this as the reduction approach. The problem whit the reduction approach is setting the maximum base stock level as VSM's management has not specified a maximum, but only that the base stock levels need to be non-negative. We could overcome this by setting the maximum base stock levels with the equal fill rate approach mentioned in Section 3.2 of the literature review.

The equal fill rate approach sets the individual item fill rates equal to the targeted fill rate of the entire inventory. For the situation under study at VSM, this would mean that we set the individual fill rates of all the 288 module types to 95%. The equal fill rate approach is commonly applied in practice due to the lack of advice on how this can be done on an individual stock-keeping unit (SKU) basis in a multi-SKU environment. The downside of the equal fill rate approach is that this method is inaccurate and leads to higher inventory values than necessary because not all SKUs have the same impact on the systems fulfilment performance, i.e. some SKUs are overstocked (Teunter et al., 2017).

The benefit of a reduction approach is that we only have to analyse a specific region of the total solution space, reducing the required computation time. However, this is also a limitation as we cannot state with certainty that this is the best performing region in the solution space. Because for example, it may be better to have low base stock levels (i.e., fill rates lower than 95%) for the modules of



machine types with low sales volumes while overstocking (i.e., fill rates higher than 95%) for the modules of machine types with high sales volumes.

In contrast to the reduction approach, the incrementing approaches we apply in this study have no restrictions regarding the regions in the solutions space they can analyse. The downside of this is the required computation time; therefore, a reduction approach might be a viable alternative as they require less computation time. However, additional research is required to determine if the region the reduction approach would analyse is the best performing region within the solution space.

5.2 Data collection

Before we can analyse the performance of the two approaches formulated in Chapter 4, we need to collect data for both approaches, i.e. sets of base stock levels. The generated sets of base stock levels differ in quality; they all can complete with 95% certainty at least 95% of machines of future demand scenarios within 10 weeks, but they all require a different investment (i.e., inventory value). Besides meeting the lead time objective, we also want to minimise the inventory value; therefore, the lower the inventory value, the better.

Since there is no feasible optimisation method that can determine the optimal set of base stock levels for the situation studied at VSM, and since we have no knowledge of the optimal or near-optimal (i.e., relaxed) set of base stock levels, we, unfortunately, have to generate a lot of possible sets in order to guarantee that the best set of base stock levels we have found is a good set. In Section 4.3.1, we established that because we want to optimise the base stock levels of 288 module types and at the same time minimise the inventory value, we need to generate at least 25,000 sets of base stock levels capable of meeting the objective to ensure that the best solution we found is a good solution.

Unfortunately, generating such a large amount of possible sets of base stock levels to guarantee that we found a good solution reduces the practical implementation of the two approaches. The GA approach requires approximately 10 minutes to generate one possible set of base stock levels, and the local search approach requires around 5 minutes for this. We consider these running times acceptable, considering that determining robust base stock levels is NP-hard. However, since we have to generate 25,000 possible sets of base stock levels for both approaches, this creates a problem. It takes the GA 174 days while continuously running to generate 25,000 sets of base stock levels and the local search approach 87 days. To improve the practical use of the two approaches, one could use the best set of base stock levels of 25,000 sets of base stock levels generated the year before as a benchmark. Then the model can, for example, be stopped as a solution has been found that is close enough to the benchmark solution.

Unfortunately, we do not have such a benchmark solution yet, and due to practical reasons, we currently do not have enough time to generate 25,000 solutions. Therefore, we continue this analysis with fewer samples (i.e., solutions) per approach. To justify the size of the data samples we use for further analysis, we analyse the spread (i.e., boxplots) of the inventory values of generated sets of base stock levels for different sample sizes. Preferable, we would compare the spread of samples to the spread of the entire population to determine from which sample size onwards the spread of the samples are comparatively similar to that of the entire population. Unfortunately, we do not know the spread of the population; as an alternative, we begin with multiple small samples of the same size and compare their spread. If the spread of the samples is not similar, we deem the sample size too small. We increase the sample size until the spread of the samples become comparable. See Appendix E for the whole sample size analysis. Based on this analysis, we conclude that from 500 sets of base stock levels per sample, they become comparatively similar to each other. Comparatively similar in this

setting means the quartile ranges of the different samples do not differ more than a few hundred thousand euros from each other, which is compared to the millions of inventory value required to achieve the objective acceptable in our opinion.

However, besides the spread of the inventory values of generated sets of base stock levels, we are also interested in the best (i.e., cheapest) set of base stock levels found per number of generated sets of base stock levels; we plotted this in Figure 5.1 for both approaches. The inventory value of the best set of base stock levels found can be read from the y-axis, while the number of generated sets is stated on the x-axis.



Figure 5.1: The best (i.e., cheapest) inventory value per number of generated sets of base stock levels for both the GA and the local search approach.

From Figure 5.1, we conclude that a sample size of 500 is large enough for the local search approach but not for the GA, as the GA finds a better solution after approximately 600 sets of base stock levels. However, it should be noted that the difference in inventory value of the cheapest solution found already considerably reduces after 35 sets of base stock levels for the GA and 175 sets of base stock levels for the LS. To be certain that the sample size of sets of base stock levels we use for further analysis of both approaches contains the characteristics of the entire population, we apply a sample size of 1,000 sets of base stock levels. Based on Figure 5.1, we also conclude that there is a very large difference in inventory value between the sets of base stock levels generated by both approaches; In Section 5.3.2, we discuss this in more depth.

Using sample sizes smaller than 25,000 comes at risk because we cannot guarantee that we found the best (i.e., most cost-effective) possible solution that we could have found for the two approaches. But by ensuring that the spread of the samples is comparatively similar, we can ensure that the most cost-effective solution we found is of the same order size as the best solution that we could have found.

5.3 Performance analysis solution approaches

The GA and the local search (LS) approach have a double objective; they need to generate sets of base stock levels capable of meeting the lead time objective with a certain degree of certainty while minimising the required inventory value. We benchmark the two approaches against each other for both these objectives. In Section 5.3.1, we focus on the performance of both approaches, i.e. their ability to ensure that with 95% certainty, at least 95% of the machines of future demand scenarios can



be completed within 10 weeks. In Section 5.3.2, we analyse the inventory value of the sets of base stock levels generated by both approaches.

5.3.1 Lead times

To ensure that each set of base stock levels can meet the objective of completing at least 95% of the machines of future demand scenarios within 10 weeks with 95% certainty, we test them by using simulation runs until the relative error of the performance is less than 5%. The performance is expressed in the percentage of machines completed within 10 weeks. Only when the average performance of the conducted simulation runs meets the objective, we store the set of base stock levels as a potential solution. If the performance is insufficient, the set is boosted by one of the approaches and retested.

By averaging the performance of the simulation runs required to ensure that the relative error of the performance of a set of base stock levels is less than 5%, we can guarantee that any generated set of base stock levels meets the objective with a certainty of 95%. To prove this, we take the best (i.e., most cost-effective) solution of both approaches and test whether they can meet the target of completing at least 95% of the machines of a one-year demand scenario within 10 weeks with a 95% degree of certainty. The inventory levels of the most cost-effective solutions can be found in Appendix F.

To determine whether the most cost-effective set of base stock levels found by both approaches can meet the lead time objective with a certainty of 95%, we apply a one-sample t-test. With the one-sample t-test, we determine if we can state with 95% certainty (i.e., confidence interval) that the performance (i.e., percentage of machines completed within 10 weeks) of the analysed set of base stock levels is at least 95% (i.e., 0.95). We analyse the performance of the best set of base stock levels found by both approaches for 250 one-year demand scenarios, expressed in the percentage of machines with a lead time below 10 weeks. We take such a large sample as we expect only small differences between the performances; therefore, we need to ensure that the dataset is large enough to get a nuanced answer. See Figure 5.2 for a graphic statistical summary report of the results of the 250 simulation runs carried out for both sets of base stock levels.



Figure 5.2: Graphic statistical summary report of the 250 executed simulation runs (i.e., scheduled one-year demand scenarios) for both approaches' best set of base stock levels.

The p-values of the Anderson Darling normality tests are less than 0.05 (i.e., 5%), and since we want 95% certainty, this means that for both approaches, we reject that the performance of the best set of base stock levels is normally distributed (i.e., null hypothesis) and instead accept that their performance is not normally distributed (i.e., alternative hypothesis). However, this does not matter because when sample sizes are large enough, t-tests are still applicable (Lumley et al., 2002).

Now that we have visualised the performance spread for both sets of base stock levels, we can conduct the one-sample t-tests to determine whether we reject or accept the null hypothesis regarding the capability of both approaches to generate sets of base stock levels capable of meeting the objective.

 $\begin{array}{ll} H_0 & \mu = \ 0.95 \ (i.e., 95\%) & (null \ hypothesis) \\ H_1 & \mu < \ 0.95 \ (i.e., 95\%) & (alternative \ hypothesis) \end{array}$

If we reject the null hypothesis, we accept the alternative hypothesis, meaning that the performance of the set of base stock levels analysed does not meet the objective of completing at least 95% of the machines of future demand scenarios within 10 weeks. We reject the null hypothesis if the p-value we determine with the one-sample t-test is below 0.05 (i.e., 5%) since we use a 95% confidence interval. See Figure 5.3 for the report of the one-sample t-tests for the most cost-effective solution that should be able to achieve the lead time objective of both approaches.

ne-Sample T: Performance GA; Performance LS						
Descriptive St	atistic	cs				
Sample	N	Mea	n StDev	SF Mean	95% Upper Bound for u	
Performance GA	250	0 9481	2 0.02341	0.00148	0 95056	
Performance LS	250	0.9482	0.02529	0.00160	0.95084	
Test	,		, ,			
Null hypothesis	ł	H₀:μ=0	,95			
Alternative hypoth	nesis H	H₁:μ<0	,95			
Sample	T-Va	lue P-	Value			
Performance GA	-1	,27	0,103			
Performance LS	-1	13	0.131			

Figure 5.3: Report of the one-sample t-test for the most cost-effective solution of both approaches.

As both sets of p-values are above 0.05 (i.e., 5%), we conclude that both sets of base stock levels do not reject the null hypothesis and thus can fulfil the lead time objective with 95% certainty. Now that we have proven that the most cost-effective set of base stock levels generated by both approaches can meet the objective with a certainty of 95%, we need to determine if one of these sets of base stock levels performs significantly better than the other. This should not be the case as both solutions are created with the same lead time objective and accuracy in mind, but in order to prove that both the GA and the local search algorithm create equally good solutions performance-wise, we conduct a two-sample t-test. With the two-sample t-test, we determine if we can state with 95% certainty (i.e., confidence interval) that the performance (i.e., percentage of machines completed within 10 weeks) of both sets of base stock levels are not significantly different from each other. See Figure 5.4 for the report of the conducted two-sample t-test.



Two-Sample T-Test and CI: Performance GA; Performance LS							
Method							
μ1: population mea μ2: population mea Difference: μ1 - μ2	n of Pe n of Pe	erformance erformance	te GA te LS				
Equal variances are	not as	sumed for	this analy:	sis.			
Descriptive Sta	atistio	s					
Sample	Ν	Mean	StDev	SE Mean			
Performance GA	250	0,9481	0,0234	0,0015			
Performance LS	250	0,9482	0,0253	0,0016			
Estimation for	D:44						
Estimation for	Diffe	erence					
Difference 959	6 CI fo	r Differe	nce				
-0,00008 (-	0,0043	6; 0,0042	0)				
Test							
Null hypothesis Alternative hypoth	esis H	Но: µл - µ; Нл: µл - µ;	₂ = 0 ₂ ≠ 0				
T-Value DF -0,04 495	P-Valu 0,9	ие 71					

Figure 5.4: Report of the two-sample t-test.

Since the p-value is above 0.05 (i.e., 5%), we do not reject the null hypothesis, which means that there is no statistically significant difference between the performance of the two sets of base stock levels, and therefore we conclude that both approaches provide equally good solutions in terms of performance.

5.3.2 Inventory value

In Section 5.3.1, we proved that the most cost-effective sets of base stock levels found by both approaches are capable of ensuring with 95% certainty that at least 95% of the machine of future demand scenarios can be completed within 10 weeks. We also proved that there was no significant difference between the performance of the two sets of base stock levels. However, we have a double objective when determining base stock levels, as besides meeting a lead time objective for a certain percentage of the machines, we also want to minimise the required inventory value. Therefore in this section, we analyse the inventory values of the sets of base stock levels generated by the two approaches.

For the analysis of the inventory values, we use 1,000 sets of base stock levels generated by each of the approaches, as discussed in Section 5.2. To visualise the spread of the inventory values of the 1,000 sets of base stock levels generated by both approaches, we plotted them against each other in Figure 5.5, sorted in ascending order. The inventory value of the sets of base stock levels can be read from the y-axis, and the number of the sets are stated on the x-axis.



Figure 5.5: The inventory values in ascending order of the 1,000 sets of base stock levels generated by the GA and local search approach.

Figure 5.5 shows a large spread in inventory value of the generated sets of base stock levels per approach, which can be explained by the fact that the sets of base stock levels are based on possible future demand scenarios. Possible future demand scenarios incorporate demand uncertainty and can therefore be considerably different from each other.

We expect the more expensive sets of base stock levels to consist of more modules than the cheaper sets of base stock levels. To test this hypothesis, we plotted the number of modules for the sets of base stock levels in Figure 5.5 in Figure 5.6 in the same sequence; for both approaches, we added a linear trendline. Based on the trendline, we conclude that more expensive sets of base stock levels contain more modules on average, confirming the hypothesis. However, it should be noted that more expensive sets of base stock levels do not necessarily consist of more modules than cheaper ones, indicating that determining the best set of base stock levels is not simply storing more modules but rather storing the right modules.





Figure 5.6: The number of modules per set of base stock levels in Figure 5.5.

We want to determine whether one of the approaches outperforms the others in generating costeffective sets of base stock levels that can meet the lead time objective. From Figure 5.5, we can already conclude based on eye-balling that the local search approach outperforms the GA approach, but we need to prove this statistically to be certain. To do so, we conduct a two-sample t-test to determine if there is a statistically significant difference between the inventory values of the sets of base stock levels generated by both approaches. For the two-sample t-test, we use the inventory value of the 1,000 sets of base stock levels generated by both approaches for this analysis. Figure 5.7 states the graphic statistical summary report of the inventory values of the 1,000 sets of base stock levels per approach.



Figure 5.7: Graphic statistical summary report of the inventory value of 1,000 generated sets of base stock levels per approach.

The p-values of the Anderson Darling normality tests are below 0.05 (i.e., 5%), and since we want 95% precision, this means that the distributions of the inventory values of the sets of base stock levels generated by both approaches are not normally distributed. We overcome this by applying large sample sizes so that t-tests are still applicable (Lumley et al., 2002).

Two-Sample T-1	Fest a	nd CI: I	nvento	ry value	s GA; Inv	ventory	values LS
Method							
μ1: population mean o μ2: population mean o Difference: μ1 - μ2	f Invento f Invento	ory values GA ory values LS	A				
Equal variances are not	assumed	d for this analy	vsis.				
Descriptive Statis	tics						
Sample	N	Mean	StDev	SE Mean			
Inventory values GA	1000	33232107	4722633	149343			
Inventory values LS	1000	26155293	2927855	92587			
Estimation for Di	fferen	ce					
959	6 CI for						
Difference Dif	ference						
7076814 (673217	0; 7421	458)					
Test							
Null hypothesis	H _o : µ ₁	- µ2 = 0					
Alternative hypothesis	Η ₁ : μ ₁	- μ₂ ≠ 0					
T-Value DF P-	Value						
40,27 1668	0,000						

Figure 5.8: Report of the two-sample t-test.

The p-value of the conducted two-sample t-test in Figure 5.8 is below 0.05 (i.e., 5%); therefore, we reject the null hypothesis meaning that there is a significant difference between the inventory values of the sets of base stock levels generated by both approaches. Looking at the spread of the 1,000 sets of base stock levels generated by both approaches in Figures 5.5 and 5.7, we conclude that the local search approach outperforms the GA significantly, proving our hypothesis.

As mentioned in Section 5.1, the commonly applied equal fill rate approach, which sets individual item fill rates equal to the target fill rate of the entire inventory, is inaccurate and leads to higher inventory values than necessary (Teunter et al., 2017). However, we still want to compare the inventory value of the set of base stock levels generated by the equal fill rate approach with the sets of base stock levels generated by the equal fill rate approach is 48.3 million euros. Based on this, we conclude that the set of base stock levels generated by the equal fill rate approach is 48.3 million euros. Based on this, we conclude that the set of base stock levels generated by the local search approach is the set of base stock levels generated by the set of base stock levels generated by the equal fill rate approach is 48.3 million euros. Based on this, we conclude that the set of base stock levels generated by the equal fill sets of base stock levels generated by the local search approach is 48.3 million euros. Based on this, we conclude that the set of base stock levels generated by the equal fill sets of base stock levels generated by the local search approach is 48.3 million euros.

Because the equal fill rate approach overstocks heavily and because all the separate fill rates are set equal to the overall fill rate, we do not really have to check if the generated set of base stock levels is capable of ensuring with 95% certainty that at least 95% of the machine of future demand scenarios can be completed within 10 weeks. But to be certain, we do; therefore, we execute 250 simulation runs with the set of base stock levels determined by the equal fill rate approach. The graphical statistical summary report of the executed simulation runs can be found in Figure 5.9.





Figure 5.9: Graphic statistical summary report of the 250 executed simulation runs for the equal fill rate approach.

Based on the statistical summary report of the executed simulation runs, we are already able to conclude that the set of base stock levels generated with the equal fill rate approach can fulfil the lead time objective. But for statistical proof, we conduct a one-sample t-test to determine if we can state with 95% certainty that the performance of the set of base stock levels meets the objective of ensuring that at least 95% of the machines of future demand scenarios can be completed within 10 weeks.

One-Sample T: Performance equal fill rate approach							
Desci	riptive S	tatistics					
				95% Upper Bound			
N	Mean	StDev	SE Mean	for μ			
250	0,96468	0,03225	0,00204	0,96805			
μ: pop	pulation me	an of C1					
Null hy	ypothesis	H ₀ :	: μ = 0,95				
Alterna	ative hypot	thesis H ₁ :	μ < 0,95				
T-Va	llue P-Va 7,20 1	alue ,000					

Figure 5.10: Report of the one-sample t-test.

As the p-value in Figure 5.10 is above 0.05 (i.e., 5%), we do not reject the null hypothesis. Therefore, we conclude that the set of base stock levels generated by the equal fill rate approach is capable of ensuring with 95% certainty that at least 95% of the machine of future demand scenarios can be completed within 10 weeks.

To get an overview, we plotted the most cost-effective sets of base stock levels generated by the GA and local search approach against the set of base stock levels generated with the equal fill rate approach in Figure 5.11. From this figure, we conclude that the local search approach is the best approach for determining cost-effective, robust base stock levels. The GA performs a bit worse but is still applicable, however the equal fill rate approach we do not deem suitable for determining cost-effective base stock levels.



Figure 5.11: Inventory value of most cost-effective sets of base stock levels generated by the GA and local search approach and the set of base stock levels generated with the equal fill rate approach.

5.4 Fulfilment of research gap

For the research gap, we identified the lack of realistic, practical approaches for determining robust base stock levels in a multi-product assemble-to-order system with complex product configurations. To fill this gap, we identified two potential approaches in Section 4.2: the GA and local search approach, from which we expected that they were capable of determining robust base stock levels for a practical, real-world scenario with multiple complex products that customers can configure order.

To determine whether one or perhaps both approaches are capable of filling the gap in research, we need to consider the extent to which they are able to do the following:

- The ability to generate robust, cost-effective sets of base stock levels capable of meeting the objective (i.e., performance);
- The ability to be applied to a full-sized real-world problem (i.e., realistic);
- The ability to take into account stochastic factors (i.e., realistic);
- The difficulty of implementation (i.e., practicality);
- The required calculation time (i.e., practicality).

In Section 5.3.1, we proved that both approaches are able to generate sets of base stock levels that can guarantee with 95% certainty that at least 95% of the machines of future demand scenarios can be completed within 10 weeks. In Section 5.3.2, we proved that the sets of base stock levels generated by both approaches are considerably cheaper than the set of base stock levels generated with the equal fill rate approach, which indicates that both approaches are a better option than the equal fill rate approach commonly applied. However, there is also a considerable difference between the sets of base stock levels generated by the GA and local search approach in terms of inventory value. The spread in inventory value of the sets of base stock levels generated by the GA. Besides that, the most cost-effective set of base



stock levels generated by the local search approach is approximately 3.5 million euros cheaper than the most cost-effective solution found by the GA. Therefore, we conclude that the local search approach performs better in terms of cost-effectiveness.

Both approaches can be applied to full-sized real-world problems, as shown by testing both approaches for the situation under study at VSM. For the case study at VSM, we consider 15 different machine types that can be configured to order by the customers; in total, 288 different module types are taken into account. By combining the GA and the local search heuristics with a simulation model of VSM's purchase, assembly and inventory management and assembly process, almost all aspects of these processes can be made stochastic to increase the simulation model's validity. The only requirement is that enough reliable data is available to determine which probability distribution best describes the process. In case there is insufficient data available, one could still apply a triangular probability distribution based on the minimum, maximum and mode estimate of the process, as we applied for the replenishment times of the modules at VSM.

The use of the simulation model in combination with either the GA or the local search enables the generation of robust, cost-effective sets of base stock levels. However, this also reduces the implementability of both approaches. Because someone who wants to apply one of these approaches would have to create a simulation model for the case study at hand, which requires knowledge about simulation modelling, reducing the implementability and thus the practicality of both approaches.

In Section 4.3.1, we have established that we need to gather at least 25,000 possible solutions to ensure that the best (i.e., most cost-effective) solution we have found is a good solution as we do not have an optimisation approach or a benchmark solution. In Section 5.1, we have established that generating 25,000 sets of base stock levels would cost the GA 174 days (i.e., approximate 10 minutes per set) and the local search approach 87 days (i.e., approximate 5 minutes per set) while continuously running. These extremely long calculation times reduce the practicality of both approaches. However, it should be noted that no code optimisation has been conducted to reduce the required running time, so there is room for improvement.

Based on the aspects discussed, we conclude that both approaches can generate robust sets of base stock levels for full-sized real-world problems but that the practical implementability of both approaches can be improved. However, we still consider both approaches a useful contribution to the literature as they help filling the research gap.

5.5 Conclusion

This chapter answers research question 4: 'What is the best approach for determining base stock levels at VSM?'. In Section 5.1, we explained that both of our approaches start with base stock levels far below the minimum required base stock levels, on which they then improve until they are able to fulfil the objective. The improvements are done randomly based on the performance of the set of base stock levels in the simulation model. Because both approaches are not exact approaches but heuristics, they cannot determine the best solution; instead, they generate a solution. Therefore, Section 5.2 determined the number of solutions we need to gather to ensure that the most cost-effective solution that can meet the objective we found is a good solution. We concluded that we need to gather 25,000 solutions; however, we cannot gather 25,000 solutions per approach due to time constraints, as this would take months. Therefore we conducted a sample size analysis and concluded that a sample size consisting of 1,000 sets of base stock levels is enough for the further analysis of both approaches.

Section 5.3 analyses the performance of the most cost-effective set of base stock levels that can meet the lead time objective generated by the GA algorithm and local search approach. We analysed their

ability to ensure with 95% certainty that at least 95% of the machines of future demand scenarios can be completed within 10 weeks. Based on one-sample t-tests, we can state with statistical certainty that both sets of base stock levels can fulfil that objective. We also conducted a paired two-sample ttest to determine if there is a significant difference between the performance of the two solutions. Based on the two-sample t-test, we concluded that there is no significant difference between the performances and that, therefore, both approaches provide equally good sets of base stock levels in terms of performance.

Besides meeting the required lead-time objective, we also want to minimise the required inventory value. Therefore, Section 5.3. also analyses the inventory values of the generated sets of base stock levels of both approaches. We used 1,000 sets of base stock levels generated by each approach for the analysis. To determine whether one of the approaches outperforms the other in generating cost-effective sets of base stock levels that can meet the lead time objective, we conducted a two-sample t-test. Based on the t-test, we determined that there is a significant difference between the two approaches. Therefore we plotted the inventory values of the 1,000 sets of base stock levels generated by each approach against each other, and based on this, we concluded that the local search approach outperforms the GA significantly.

For completeness, we also compared the inventory values of the sets of base stock levels generated by the GA and local search approach to the inventory value of the set of base stock levels generated with the commonly applied equal fill rate approach. Based on this comparison, we concluded that the set of base stock levels generated by this approach is outperformed in terms of cost-effectiveness by all sets of base stock levels generated by the local search approach and most of the sets of base stock levels generated by the GA approach.

Based on the conducted analyses, we conclude that the local search approach is best for determining cost-effective, robust base stock levels that can meet the objective of ensuring with 95% certainty that at least 95% of the machine of future demand scenarios can be completed within 10 weeks. The GA performs a bit worse but is still applicable. We do not recommend applying the equal fill rate approach as we proved this leads to considerable overstocking.

In Section 5.4, we discussed whether the GA and local search approach are a valuable contribution to the research gap. The aspects based on which we assess the approaches are:

- The ability to generate robust, cost-effective sets of base stock levels capable of meeting the objective (i.e., performance);
- The ability to be applied to a full-sized real-world problem (i.e., realistic);
- The ability to take into account stochastic factors (i.e., realistic);
- The difficulty of implementation (i.e., practicality);
- The required calculation time (i.e., practicality).

Based on the assessment of both approaches, we concluded that they are both suitable for generating robust sets of base stock levels for full-sized real-world problems but that the practical implementability of both approaches can be improved. Nonetheless, we consider both approaches a useful contribution to the literature.


6. Conclusions and discussion

In Chapter 5, we analysed the suitability of the GA and local search approach we developed for determining robust and cost-effective sets of base stock levels for the situation under study at VSM. This chapter concludes and discusses the research we conducted to answer the main research question. Section 6.1 states the main conclusion to this research, and Section 6.2 contains the recommendations we have for VSM regarding the implementation of the developed approaches. Section 6.3 covers the limitations of the developed approaches and possible future research we discuss in Section 6.4.

6.1 Conclusion

This research aimed to develop a method for determining the required base stock levels to meet a 10week lead time objective for at least 95% of the machines of future demand scenarios with 95% certainty while minimising the inventory value. The main research question we formulated to achieve this goal is:

How to determine which modules should be kept in stock and in what quantity so that VSM can guarantee a 10-week lead time for at least 95% of the orders while minimising inventory value?

To answer this question, we applied the managerial problem-solving method of Heerkens and van Winden (2012) and divided the solution process into 5 stages being: (1) current situation, (2) literature review, (3) solution design, (4) analysis of results, and (5) conclusion and recommendations. This section analyses the findings for these stages.

Currently, VSM applies a discrete order quantity (DOQ) inventory management method. DOQ means that parts are purchased in the exact amount needed as VSM does not want to keep free to use anonymous inventory due to the enormous amount of different parts it uses. DOQ inventory management methods can only be fully utilized in case of neglectable ordering costs and replenishment times (Gosrani and Kolekar, 2017). Unfortunately for VSM, the replenishment times are not neglectable and, in fact, they have a major impact on the final lead times of the machines. To overcome this, VSM started forecasting expected sales and machine configurations to be able to order parts upfront. Unfortunately, the forecasting precision leaves much to be desired, causing the wrong parts to be ordered.

Consequently, the machines still do not have the desired lead times, and about 7 million euros worth of wrongly ordered parts are laying in storage. This negates the reasons for choosing the DOQ strategy for inventory management. To overcome this problem, we opted to store modules as we think these can act as a buffer against the negative impact of part lead times on machine lead times, eliminating the need to forecast expected sales (i.e., orders are not planned until they are actually placed). To determine the validity of this theory, VSM's management wanted to have a method capable of accurately determining which modules need to be kept in stock and in what quantity. The storage of parts would also be a viable option to eliminate the negative impact of part lead times on machine lead times on machine lead times. In this study, we opted for modules, as it proved difficult to determine the parts required for the modules accurately.

To gain insight into aspects that are normally considered by assembly strategies for modular products and in already developed methods for determining base stock levels, we studied the academic literature. Based on Olhager (2010), we concluded that VSM applies an assemble to order (ATO) strategy, and from Vos (2001), we learned that the successful realisation of a flexible ATO requires

suitable solutions to three key problems: (1) product design, (2) assembly layout design, and (3) planning method.

Good modular product design means that several modules can be paired to the same base product to create products with different characteristics (Ulrich, 1994). Based on the conclusion of Hölttä-Otto and Weck (2007) that determining the most appropriate product modules is difficult as it requires indepth technical knowledge of the product, we decided to use the machine modules currently used by VSM in this study. Based on the conclusion of Battini et al. (2007) that fixed-position assembly layouts with a one-piece flow are the best option for low production volumes with high complexity and long set-up times, we have decided that the currently applied assembly layout of VSM is perfectly suitable for their situation. For the planning method, we concluded together with VSM's management based on the description of the priority rule-based heuristic given by De Boer (1998) that this would best represent the applied scheduling method.

After reviewing the literature for already developed methods to determine base stock levels, we concluded based on Teunter et al. (2017) that simply equating the SKUs' item fill rates to the target leads to non-cost-effective inventory levels. Based on Thonemann et al. (2002), Agrawal and Cohen (2001), Daniel and Rajendran (2005), Avsar et al. (2009), Albrecht (2014) and Song and Yao (2002), we concluded that several exact, heuristic, queuing, and simulation methods have been developed over the years but that due to the complexity of the problem they are either highly theoretical or simplified. Therefore, we decided that developing a realistic, practical approach suitable for multi-product assemble-to-order situations with complex product configurations is necessary to answer the main research question.

Consequently, the next step was to decide which approach would be most suitable to achieve this. We concluded that conducting experiments with VSM's real assembly system was not practical. We also decided that the enumeration of all possible solutions is not realistic because of the huge solution space. Since Bienstock and Özbay (2008) proved that determining robust base stock levels is NP-hard, we unfortunately also had to conclude that exact optimisation methods cannot be applied. Therefore, we decided that applying heuristic approaches would be the most promising. We came up with two heuristic approaches that we considered the most viable.

The first approach is a genetic algorithm (GA); we concluded that GAs might be suitable for determining robust base stock levels because of their ability to deal with real-life size problems and their ability to use historical data to guide the search to the best performing region within the solution space (Daniel and Rajendran, 2005). The second approach is a combination of a GA and local search algorithm; we refer to this approach as the local search approach. We opted for this combination as GAs are not well suited for fine-tuning solutions (i.e., local search), which are very close to optimal ones (Martinez and Lozano, 2007). The GA guides the search to the best performing region within the solution space, after which the local search takes over to analyse this region. We selected local search as it has proven to be very successful in determining near-optimal and sometimes even optimal solutions for difficult reallife problems with enormous solution spaces (Aarts and Lenstra, 2003; Dumitrescu and Stützle, 2003). As both approaches are heuristics and not exact algorithms; they cannot guarantee that they found the optimal solution; instead, they can generate sets of base stock levels that are capable of meeting the objective of completing at least 95% of the machines of future demand scenarios within 10 weeks with 95% certainty. Therefore, we let the heuristics generate multiple sets of base stock levels and select the most cost-effective set of base stock levels that can meet the objective as the solution. To ensure that the selected most cost-effective set of base stock levels is a good solution, we need to gather at least 25,000 solutions. The number of 25,000 is based on the number of module types we



take into account and the dual objective of minimising the inventory value while meeting a lead time objective.

Because we decided that conducting experiments with VSM's real assembly system is not practical, we had to find another method to determine the performance of the sets of base stock levels generated by the GA and local search approach. We concluded that the best approach was to create a simulation model of VSM's purchase, inventory management, and assembly process. The simulation model had to be able to determine what machine lead times can be realised for a given demand scenario and set of base stock levels.

With the help of the simulation model, we analysed the performance of the GA and local search approach. We analysed if the sets of base stock levels generated by the approaches were capable of ensuring with 95% certainty that at least 95% of the machines of future demand scenarios can be completed within 10 weeks. We did this by running multiple future demand scenarios with the sets of base stock levels to see how they performed; by conducting some one-sample t-tests based on the performance and the objective, we conclude with statistical certainty that both approaches can create sets of base stock levels that can fulfil the lead time objective. However, besides meeting the required lead-time objective, we also want to minimise the required inventory value. Therefore, we analysed 1,000 sets of base stock levels generated by each approach. We conducted a two-sample t-test and concluded that the local search approach creates significantly cheaper solutions (i.e., sets of base stock levels).

We also compared the inventory values of the sets of base stock levels generated by the GA and local search approach to the inventory value of the set of base stock levels generated with the commonly applied equal fill rate approach. Based on this comparison, we concluded that the set of base stock levels generated by this approach is outperformed in terms of cost-effectiveness by all sets of base stock levels generated by the local search approach and most of the sets of base stock levels generated by the GA and local search approach against the set of base stock levels generated with the equal fill rate approach in Figure 6.1.



Figure 6.1: Inventory value of most cost-effective sets of base stock levels generated by the GA and local search approach and the set of base stock levels generated with the equal fill rate approach.

Based on the conducted analyses, we conclude that the local search approach is best for determining cost-effective, robust base stock levels that can meet the objective of ensuring with 95% certainty that at least 95% of the machines of future demand scenarios can be completed within 10 weeks. The GA performs a bit worse but is still applicable. We do not recommend applying the equal fill rate approach as we proved this leads to considerable overstocking.

The last step was to determine if the approaches we developed and tested with the situation under study at VSM contribute to the literature and help fill the research gap. Based on the assessment of both approaches, we concluded that they are both suitable for generating robust sets of base stock levels for full-sized real-world problems but that the practical implementability of both approaches can be improved. Nonetheless, we consider both approaches a useful contribution to the literature.

6.2 Recommendations

We proved that the local search approach performs best for determining the most cost-effective set of base stock levels that, with 95% certainty, can guarantee that at least 95% of the machines of future demand scenarios can be completed within 10 weeks. However, if this approach really generates the base stock levels needed for the real world situation at VSM depends heavily on the validity of the applied simulation model. This is the major weak spot of the applied approach, and therefore we strongly recommend that VSM's management gathers the data needed to validate the simulation model further.

To improve the model's validity, we recommend that VSM gathers a year's worth of data, including demand, machine configurations, order dates, realised lead times, module replenishment times and which modules are incorrectly forecasted. With the gathered data, a real demand scenario can be created based on which the performance of the simulation model measured in machine lead times can be compared to the actually realised machine lead times. For this, the simulation model needs to be altered a bit to be able to take into account wrongly forecasted machine configurations. Such an analysis can determine if the used priority rule-based scheduling heuristic really represents VSM's scheduling process. It also enables further validation of other aspects of the simulation model, such as the purchasing (i.e., replenishment) and inventory management processes.

If the simulation model is deemed realistic after validating, we recommend that VSM conducts additional research to speed up the applied local search algorithm. This can be done by changing the operational procedure of the machine assembly scheduling process to discrete event-based, meaning that the scheduling heuristic jumps from event to event. Events can be the arrival of parts or the completion of an assembly, testing or loading task. All events occur at a specific time that marks a change of state in the system, allowing the planning system to efficiently move from one event to the other instead of having to check in every simulation week, as is currently the case, whether the parts for an outstanding order have arrived or whether an assembly, testing or loading task has been completed or can be started. Besides altering the code for the scheduling heuristic, we also recommend that VSM conducts additional research on other aspects of the code to determine if their computation time can be reduced. In the unfortunate case that the model is not deemed realistic, we recommend that VSM first invests time into improving the model's validity.

Once the validity of the simulation models is proven and the code is optimised, we recommend that VSM's management let the local search algorithm generate 25,000 possible sets of base stock levels. After generating the 25,000 sets, we recommend that VSM's management studies the most cost-



effective set of base stock levels with great care to determine if the base stock levels generated by the model make sense based on their personal insight of the real assembly process. In the unfortunate case that the found base stock levels seem strange, we recommend that VSM tries to determine what could have caused this and, if needed, improves the simulation model based on the found insights.

The frequency whit which we recommend VSM's management to create 25,000 sets of base stock levels depends heavily on the time needed to create them. With the current computation time of the local search algorithm, we recommend an annual use, but if the computation time can be reduced considerably, VSM can also apply it more frequently in a sort of rolling forecast application per quartile. However, it should be noted that this is only beneficial in case of changing demand patterns as otherwise changing the base stock levels has no effect on the machine lead times but only creates additional work. However, it can be applied to check the applied base stock levels more frequently.

If the most cost-effective set of base stock levels makes sense, we recommend that VSM's management abolishes the current forecasting approach and implements the module base stock level assembly strategy. We recommend gradual implementation by either machine type or product segment (e.g., beam, plate, flat and angle, surface) rather than for all machine types at once to see if unexpected events occur. Interesting to note is that the module base stock level assembly strategy also incorporates forecasting through the possible future demand scenarios generated by the simulation model on which the sets of base stock levels are composed.

6.3 Limitations

As already explained in Section 6.2 regarding recommendations to VSM's management, the validity of the simulation model is the biggest limitation of the applied method. Other limitations are that although this research aimed to develop a practical, realistic method for determining base stock levels given a certain objective, we had to simplify the method on some aspects. The aspects simplified in relation to the real world are the storage of modules and the assembly and testing of the machines and modules.

The simplifications applied regarding the storage of the modules are that we do not consider the required storage space and inventory holding cost when determining the base stock levels, whereas these can have a rather large impact on the operational costs. We have left these out of consideration as VSM's management has not yet determined how to stock the pre-assembled modules.

The simplifications applied regarding the assembly and testing of the modules and machines are that the required assembly and test duration are deterministic while these have a stochastic character depending on, for example, the experience of the employees. The reason for using deterministic durations is that no data to determine the stochastic durations are available. Another simplification regarding the assembly and testing is that testers can test all machine types, and mechanics can assemble all machines types. In reality, this is not the case due to the machines' wide variety and complexity. One more simplification is that we use a constant test and assembly capacity (i.e., hours) during the year, not taking into account holidays or sickness of the employees.

Another limitation of the model is that the stochastic customer demand generation is based on historic demand, while historic demand patterns might differ from future demand patterns due to changed settings such as phasing out machines or introducing new machines. Also, ordering costs are not taken into account due to the applied DOQ inventory management strategy.

6.4 Future research

This research proved that the local search approach is suitable for determining cost-effective sets of basic stock levels. However, we have not analysed the generated sets of base stock levels themself in detail, which leaves a blind spot for which we recommend additional research. We want to get a detailed insight into the height of the base stock levels per module type. We would like to know whether certain module types always have higher or lower base stock levels than others and why this is the case. Based on the insight, the intelligence of the local search approach can be improved, which we believe can significantly improve the quality of the generated sets of base stock levels in terms of performance and inventory value. We also recommend additional research regarding the boosting process to gain more insight into its performance. For this, we recommend a probing approach to gain insight into the boosting choices made, based on which we can determine whether the boosting process can be improved.

We also strongly recommend future research regarding other approaches that, together with the simulation model, can generate robust sets of base stock levels for the situation under study. Such research can prove that the local search approach is indeed the most suitable or that another approach might be better.

In this study, we used the simulation model of VSM's purchase, inventory management and assembly process as a digital twin of the real system to test the performance of the sets of base stock levels. However, another angle for further research could be using the simulation model to gain insight into other operational aspects, such as determining the most optimal assembly scheduling method, parts ordering strategy, or storage and order pick approach. However, it should be noted that the simulation model first needs to be validated more to do this.

Also interesting would be turning the objective around and developing a method for determining the best base stock levels for a given budget. Such a method would provide managers with insight into the machine lead times that can be achieved with a certain budget, which is something that the approaches we analysed in this study are lacking.



References

Aarts, E. H., & Lenstra, J. K. (2003). Local search in combinatorial optimization. Princeton University Press.

Agrawal, N., & Cohen, M. A. (2001). Optimal material control in an assembly system with component commonality. *Naval Research Logistics (NRL), 48*(5), 409-429.

Albrecht, M. (2014). Determining near optimal base-stock levels in two-stage general inventory systems. *European Journal of Operational Research*, *232*(2), 342-349.

Asadi, N., Jackson, M., Fundin A. (2019). Implications of realising mix flexibility in assembly systems for product modularity—A case study. Journal of Manufacturing Systems, 52, 13-22.

Avsar, Z. M., Zijm, W. H., & Rodoplu, U. (2009). An approximate model for base-stock-controlled assembly systems. IIE transactions, 41(3), 260-274.

Axsäter, S. (1993). Continuous review policies for multi-level inventory systems with stochastic demand. Handbooks in operations research and management science, *4*, 175-197.

Battini, D., Faccio, M., Ferrari, E. (2007). Design configuration for a mixed-model assembly system in case of low product demand. International Journal of Advanced Manufacturing Technology, 34, 188–200.

Bienstock, D., & Özbay, N. (2008). Computing robust basestock levels. *Discrete Optimization*, *5*(2), 389-414.

Blum, C., & Roli, A. (2003). Metaheuristics in combinatorial optimization: Overview and conceptual comparison. ACM computing surveys (CSUR), 35(3), 268-308.

Boysen, N., Fliedner, M., Scholl, A. (2009). Sequencing mixed-model assembly lines: Survey, classification and model critique. European Journal of Operational Research, 192(2), 349-373.

Bozarth, C., & Chapman, S. (1996). A contingency view of time-based competition for manufacturers. International Journal of Operations & Production Management, 16(6), 56-67.

Cardós, M., Babiloni, E., Palmer Gato, M. E., Guijarro, E. (2013). A heuristic to minimise the inventory value of repairable parts with service constraints: Application to an airline company. Journal of Industrial Engineering and Management (JIEM), 6(3), 771-778.

Closs, D. J., Nyaga, G. N., Voss, M. D. (2010). The differential impact of product complexity, inventory level, and configuration capacity on unit and order fill rate performance. Journal of Operations Management, 28(1), 47-57.

Bani-Hani, D. (2020). Genetic Algorithm (GA): A Simple and Intuitive Guide. Towards Data Science. https://towardsdatascience.com/genetic-algorithm-a-simple-and-intuitive-guide-51c04cc1f9ed

Daniel, J. S. R., & Rajendran, C. (2005). A simulation-based genetic algorithm for inventory optimization in a serial supply chain. International Transactions in Operational Research, 12(1), 101-127.

De Boer, R. (1998). Resource-constrained multi-project management (Doctoral dissertation, PhD thesis, University of Twente, The Netherlands).

Donk, P.D., & Doorne, R. (2016). The impact of the customer order decoupling point on type and level of supply chain integration. International Journal of Production Research, 54(9), 2572-2584.

Dumitrescu, I., & Stützle, T. (2003). Combinations of local search and exact algorithms. In Workshops on Applications of Evolutionary Computation. Springer, Berlin, Heidelberg.

Eiben, A. E., & Smith, J. E. (2003). Introduction to evolutionary computing, Berlin: springer, 53(18).

Fujimoto, T., & Nobeoka, K. (2004). In which industries does Japan excel? The compatibility between architecture and organizational capability, Research Institute of Economy, Trade and Industry(RIETI).

Gaither, N., & Frazier, G. (2004). Operations Management (9th ed.). Cengage Learning.

García-Martínez, C., & Lozano, M. (2007). Local search based on genetic algorithms. In Advances in metaheuristics for hard optimization. Springer, Berlin, Heidelberg.

Gayam, N.R., Shanmuganandam, K., Vinodh, D. (2020). Layouts in production industries: A review. Materials Today: Proceedings.

Gen, M., Cheng, R. (2000). Genetic Algorithms and Engineering Optimisation. Wiley, New York.

Gosrani, V., & Kolekar, A. (2017). MRP using DOQ and EOQ lot Sizing Techniques. International Journal Innovative Research Science English Technology, 6, 4110-4117.

Hans, E.W., Herroelen, W., Leus, R., Wullink, G. (2007). A hierarchical approach to multi-project planning under uncertainty. Omega, 35(5), 563-577.

He, D.W., & Kusiak, A. (1997). Design of assembly systems for modular products. IEEE Transactions on Robotics and Automation, 13(5), 646-655.

Heerkens, J. M. G., & van Winden, A. (2012). Geen probleem, een aanpak voor alle bedrijfskundige vragen en mysteries. Business School Nederland.

Heike, G., Ramulu, M., Sorenson, E., Shanahan, P., & Moinzadeh, K. (2001). Mixed model assembly alternatives for low-volume manufacturing: the case of the aerospace industry. International Journal of Production Economics, 72(2), 103-120.

Herroelen, W., & Leus, R. (2005). Project scheduling under uncertainty: Survey and research potentials. European journal of operational research, 165(2), 289-306.

Hölttä-Otto, K., & de Weck, O. (2007). Degree of Modularity in Engineering Systems and Products with Technical and Business Constraints. Concurrent Engineering, 15(2), 113-126.

Hu, K. C., Tsai, C. W., & Chiang, M. C. (2020). A multiple-search multi-start framework for metaheuristics for clustering problems. IEEE Access, 8, 96173-96183.

Koren, Y., Heisel, U., Jovane, F., Moriwaki, T., Pritschow G., Ulsoy G., Van Brussel H. (1999). Reconfigurable Manufacturing Systems. CIRP Annals, 48(2), 527-540.

Law, A.M (2015). Simulation Modeling and Analysis (5th e.d.). McGraw-Hill.

Lumley, T., Diehr, P., Emerson, S., Chen, L. (2002). The importance of the normality assumption in large public health data sets. Annual review of public health, 23(1), 151-169.

Maalouf, E., Le Duigou, J., Hussein, B., Daaboul, J. (2020). Distributed Scheduling in Cellular Assembly for Mass Customization. IFIP International Conference on Product Lifecycle Management, 3-14.

Mourtzis, D. (2016). Challenges and future perspectives for the life cycle of manufacturing networks in the mass customisation era. Logistics Research, 9(1), 1-20.



Olhager, J. (2003). Strategic positioning of the order penetration point. International Journal of Production Economics, 85(3), 319–329.

Olhager, J. (2010). The role of the customer order decoupling point in production and supply chain management. Computers in industry, 61(9), 863-868.

Paes, I.F., Sabioni, R.C., Daaboul, J., Le Duigou, J., (2018). Product and process modular design: a review. 8th International Conference on Mass Customization & Personalization - Community of Europe (MCP-CE 2018) - Digital Customer Experience, Novi Sad, Serbia.

Pourtalebi, S., Horváth, I., Opiyo, E. (2013, August). Multi-aspect study of mass customization in the context of cyber-physical consumer durables. Conference, International Design Engineering Technical Conferences and Computers and Information in Engineering Conference, (Vol. 55911, p. V004T05A006), Portland, Oregon, USA.

Rader, D. J. (2010). Deterministic operations research: models and methods in linear optimisation. John Wiley & Sons.

Rahimi-Vahed, A., & Mirzaei, A. H. (2007). A hybrid multi-objective shuffled frog-leaping algorithm for a mixed-model assembly line sequencing problem. Computers & Industrial Engineering, 53(4), 642-666.

Rothlauf, F. (2011). Design of modern heuristics: principles and application. Springer Science & Business Media.

Sarker, B.R., & Pan, H. (2001). Designing a mixed-model, open-station assembly line using mixed-integer programming. Journal of Operational Research Society, 52, 545–558.

Schilling, M.A. (2000). Toward a general modular systems theory and its application to interfirm product modularity. Academic of Management Review, 25(2), 312-334.

Shaik, A.M., Rao, V.V.S.K., Rao, C.S. (2015). Development of modular manufacturing systems—a review. The International Journal of Advanced Manufacturing Technology, 76, 789–802.

Sharman, G. (1984). The rediscovery of logistics. Harvard Business Review, 62(5), 71-80.

Shoval, S., Efatmaneshnik, M., Ryan, M. J. (2017). Probabilistic approach to modular assembly. IFAC-PapersOnLine, 50(1), 5688-5693.

Silver, E. (2004). An overview of heuristic solution methods. Journal of the Operational Research Society, 55, 936–956.

Sivanandam, S. N., & Deepa, S. N. (2008). Genetic algorithms. Introduction to genetic algorithms, 15-37. Springer, Berlin, Heidelberg.

Song, J. S. (1998). On the order fill rate in a multi-item, base-stock inventory system. Operations Research, 46(6), 831-845.

Song, J. S., & Yao, D. D. (2002). Performance analysis and optimization of assemble-to-order systems with random lead times. Operations Research, 50(5), 889-903.

Stadtler, H., Kilger, C., Meyr, H. (2015). Supply chain management and advanced planning: concepts, models, software, and case studies.

Stevenson, W. J., & Hojati, M. (2007). Operations management. Boston: McMcGraw-Hill Irwin.

Teunter, R. H., Syntetos, A. A., Babai, M. Z. (2017). Stock keeping unit fill rate specification. European Journal of Operational Research, 259(3), 917-925.

Thengvall, B. (2019). How Many Simulation Trials Should I Run? Anylogic. https://www.anylogic.com/blog/how-many-simulation-trials-should-i-run-optquest/

Thonemann, U. W., Brown, A. O., Hausman, W. H. (2002). Easy quantification of improved spare parts inventory policies. Management Science, 48(9), 1213-1225.

Ulrich, K. (1994). Fundamentals of product modularity. In Management of Design, 219-231

Voortman Steel Group. (2021). Voortman Steel Group – Voortman Steel Group <u>https://www.voortmansteelgroup.com/en/about-voortman</u>

Voortman Steel Group. (2021). Voortman Steel Group - Voortman Steel Machinery <u>https://www.voortmansteelgroup.com/en/voortman-steel-machinery</u>

Vos, J.A.W.M. (2001). Module and system design in flexible automated assembly. Delft University Press.

Whitley, D. (1994). A genetic algorithm tutorial. Statistics and computing, 4(2), 65-85.

Wikner, J., & Rudberg, M. (2005). Integrating production and engineering perspectives on the customer order decoupling point. International Journal of Operations & Production Management, 25(7), 623-641.



Appendix

Appendix A	VSM's product range
Appendix B	Forecasting accuracy
Appendix C	Modules of the considered VSM machine types
Appendix D	Determined Poisson distribution and seasonality per machine type
Appendix E	Sample size analysis
Appendix F	Base stock levels of the best solutions

A. VSM's product range

- A1. Beam processing
- A2. Plate processing
- A3. Flat and angle processing
- A4. Surface treatment
- A5. Back to Back & Split systems



Figure A.1: All machining processes that the VSM machines can perform.





VB RANGE SAWING

The Voortman VB sawing machines are heavy, robust profile sawing machines. As fully automated machines, they can run overnight without an operator, as such improving the overall productivity of the factory.

- Ø 45°/60° miter cutting
- Ø Fast material approach
- Automatic saw band pressure adjustment
 Short piece removal system (SPRS)
- Combine with any drill for back-to-back system



Figure A.2: Machine information of the saw VB1050 and the VB1250.





Figure A.4: Machine information of the V613.





Figure A.5: Machine information of the V630-1050.



VOORTMAN.NET/v630



V630

- Simultaneous processing on all sides
- Able to process variety of profiles/sizes
- Field-proven and stable design
- Robust machine
- O Low maintenance

Customers who require the highest productivity and throughput can turn to the Voortman V630 CNC drilling machine. Lower labor costs with an automated infeed, outfeed, and automatic toolchanger for each of the 3 drill spindles. The V630 is perfect for fully-automated beam lines.



VOORTMAN.NET/v830



V630

- Simultaneous processing on all sides
- Able to process variety of profiles/sizes
- Field-proven and stable design
- Robust machine
- O Low maintenance

Customers who require the highest productivity and throughput can turn to the Voortman V630 CNC drilling machine. Lower labor costs with an automated infeed, outfeed, and automatic toolchanger for each of the 3 drill spindles. The V630 is perfect for fully-automated beam lines.



Figure A.6: Machine information of the V630-1250.



Figure A.7: Machine information of the V631-1050.



V631 DRILLING AND MILLING

- Ø Fastest milling in the market
- High autonomy level
- Direct drive motors
- Multiple functions combined
- Improved user-friendliness

The Voortman V631 adds increased milling capabilities to its proven automated beam drilling machines. Capable of simultaneous drilling, marking and milling on all sides, the V631 provides high speed processing to accelerate your production performance.





Figure A.8: Machine information of the V631-1250.





Figure A.9: Machine information of the V704.

VOORTMAN.NET/v704

V704 MARKING



- Marking by milling on all four sides
- Ø Marks visible after blasting and painting
- Full or partial contour marking
- Runs without operator
- High precision linear guides

In many operations, marking is a bottleneck. Leave the days of layout marking bottlenecks in the past with the Voortman V7D4 beam marking machine. Dnce integrated into a fully automated production line for unmanned operation, overall productivity increases.





Figure A.10: Machine information of the V807.

V807

ROBOTIC COPING / THERMAL CUTTING

- Bolt holes, copes, slots, weld preps, notches, bevels etc.
- Ø 7-axis robotic thermal cutting cell
- Very small footprint; multiple machines in one
- Sully automated processing from loading to unloading
- High definition plasma and oxy-fuel cutting

Processes like drilling, milling, sawing and acetylene cutting can either restrict the range of operations or be quite time-consuming compared to robotic plasma processing. In many cases the V807 can be not only an efficient replacement for these processes combined, but it can provide even more possibilities for future growth while increasing the speed of processing.



V808

Copes, bolt holes, slots, notches, bevels, marking – the Voortman V808 can process it while reaching all sides of the material with exceptional accuracy and speed. Both plasma and oxy-fuel torches are available. The V808 opens the door to a significant increase in production speed and flexibility.

- High definition plasma and oxy-fuel cutting
- 8-axis robotic thermal cutting cell
- Small footprint; multiple machines in one
- Improved processing speed and flexibility
- Ø Bolt holes, copes, slots, notches, cut-to-length









VOORTMAN.NET/v2000



V2000

CAMBERING

- Section bending
- Beam straightening
- Stroke memory for repetition
- Hydraulic cylinder with pusher
- Motor driven, adjustable rolls

The Voortman V2000 cambering machine is designed to camber and straighten beams. A rigid steel C-frame, horizontal cylinder pusher and memory settings, generate a repeatable and reliable output every time, all controlled by an easy-touse touchscreen operator panel.



Operating width (min)Operating width (mex)50 mmDerating width (mex)440 tonnesWorking height850 mm850 mm

Figure A.12: Machine information of the V2000.

A.2 Plate processing



Figure A.13: Machine information of the V200.







Figure A.14: Machine information of the V302.

VOORTMAN.NET/V302



PLATE CUTTING

V302

- Reduced cut-to-cut cycle time
- Fast height control
- Xtensive Hole Performance for quality holes
- Easy remounting of torch after collision
- Revolutionary oxy-fuel torch

The Voortman V302 is an effective plate cutting machine that gives smaller workshops all of the benefits from Voortman's renowned cutting technologies and reputable build quality. It is designed to cut steel plates with plasma or oxyfuel in a fast and efficient way.





VOORTMAN.NET/v302



Figure A.15: Machine information of the V303.

PLASMA AND OXY CUTTING (MOVING GANT

- O More interaction & feedback to make things easy
- O Perfect synergy between machine and workshop
- Next step in automatic bevel cutting with Xtensive Bevel Technology
- Appoint operators as workshop managers
- O Maximum up-time and capacity

V303

All new functions on the V303 contributing to more interaction and feedback, align very well with existing SigmaNEST modules. Using SigmaNEST advanced nesting solution in combination with Voortman machines, allows you to optimize your whole fabrication process from work preparation to end product thus improving your business profitability. Let's dive somewhat deeper into the specific SigmaNEST modules and the benefits you gain from it.



VOORTMAN.NET/v304



V304 PLATE CUTTING

- .
- Low X-rail for easy loading & unloading
- Reduced cut-to-cut cycle time
- Fast height control
- ⊘ I-cut correction for reducing plasma taper
- Oross inhibitor for reducing piercing dross

The Voortman V304 plate cutting machine is designed for the next step in plasma and oxy-fuel cutting quality, enabling you to cut plates in a fast and efficient way. Multiple material types can be processed such as mild steel, stainless steel and aluminum.



Figure A.16: Machine information of the V304.





Figure A.18: Machine information of the V320.





Figure A.19: Machine information of the V325.

A.3 Flat and angle processing



Flet steel (min)

Angle steel (min)

V550-7 PUNCHING, SHEARING AND DRILLING

The Voortman V550-7 flat and angle processing machine is built for top speed, reliability and versatility in punching, shearing and drilling flat and angle profiles. Do it all with multiple process capabilities included, like thread tapping, countersinking and miter shearing up to 45°.

- Fastest machine in the market
- Automatic die and punch selection
- Hydraulic numbering
- Automatic infeed buffering
- Optional automatic outfeed sorting system



Vertical punchesPunch capacity7 pcs110 tonnesHorizontel punchesPunch capacity1 pc70 tonnes

Flat steel (max)

Angle steel (max)

5x50 500x25

5x50 200x16

Figure A.20: Machine information of the V550-7.











A.4 Surface treatment



VSB RANGE

SHOT BLASTING

Voortman shot blasting machines are designed using the highest quality components for shot blasting plates and profiles. Pre-installed blasting programs and automatic functions make the VSB Range fast and easy to work with. Its compact design reduces the footprint without compromising functionality and quality. Maximize footprint efficiency by combining with painting, pre-heating, drying and a blow-off unit.

- Runs without operator
- Automatic profile batching
- Automatic shot angle adjustments
- O Long component life
- Easy maintenance



Figure A.24: Machine information of the VSB.



Figure A.25: Machine information of the VP.



A5. Back to Back & Split systems



Figure A.26: Back-to-back (Drill Saw combined) and split systems (MSI).

B. Histogram forecasting accuracy

Figure B.1: Histogram of the accuracy of the forecasting method applied by VSM.

Confidential

C. Modules of the considered VSM machine types

Table C.1: Table with the modules per machine type and their value, deterministic assembly time in hours and minimum, maximum and mode replenishment time in weeks.

Confidential

D. Poisson distribution and seasonality per machine type

We found that the Poisson distribution fits well with the real demand pattern for VSM machines. We established this by counting the demand per machine type per month; we analysed the sales data from 2015 to 2020. Based on this, we determined the average demand per month, as this is the Lambda we need as the input parameter for the Poisson formula. The Poisson formula is shown below:

$$p(x) = \frac{e^{-\lambda_{*\lambda}x}}{x!} \quad \substack{\lambda = \text{ average demand} \\ x = \text{ number of machines sold}}$$

From the Poisson distribution, we can determine the probability of the height of the demand in a certain month for each machine type. Due to this, we can sample the expected demand per month by applying inverse transform sampling using the following formula:

Demand per month =
$$\sum_{x=1}^{X} \left(U(0,1) \le \frac{e^{-\lambda_{*\lambda}x}}{x!} \right) * x$$

Now that we know the demand, we need to determine the order week, i.e. seasonality. We start by taking the square root of the number of observations, i.e. historic sales, to determine the number of seasons within a year for each machine type. We then divide the 52 weeks of a year over these seasons and determine which historic sales percentage falls within each bucket. Now we know for every machine type the seasons and the probability that a machine can be ordered in a specific season. However, the season covers multiple weeks, and we want to assign the machines an order week. Therefore, when a season is selected, we sample from the weeks within the season, the probability of the weeks is uniform.

The fit between the Poisson distribution and the observed demand patterns is given per machine type in the subsections below. Also stated per machine type are the seasons and their probabilities.

Confidential

E. Sample size analysis

The spread of the analysed sample sizes are evaluated by looking at their box plots quartile ranges (i.e., maximum, quartile 3, median, quartile 1, minimum); see Figure E.1 and the variance per quartile range.





Figure E.1: Boxplot quartiles.

Sample size 250

Table E.1: Boxplot quartiles and their variance for four samples with a sample size of 250 sets of base stock levels generated by the GA approach.

Confidential

Table E.2: Boxplot quartiles and their variance for four samples with a sample size of 250 sets of base stock levels generated by the local search approach.

Confidential

For both the GA and LS approach, we analysed 4 samples with 250 sets of base stock levels per sample; we deemed the variances to be large, especially for the GA. Therefore, we increased the sample sizes to 500 sets of base stock levels per sample.

Sample size 500 Table E.3: Boxplot quartiles and their variance for two samples with a sample size of 500 sets of base stock levels generated by the GA approach. Confidential

Table E.4: Boxplot quartiles and their variance for two samples with a sample size of 500 sets of base stock levels generated by the LS approach.

Confidential

We analysed 2 samples with 500 sets of base stock levels per sample; we deemed the variances to be acceptable as they are only a few hundred thousand euros at maximum, which, compared to the inventory values of several million, is acceptable. To increase the certainty that the sample sizes used for the analyses of the GA and local search approach are representable of the entire population (i.e., 25,000 sets of base stock levels), we double the sample size.

Sample size 1,000

We deem sample sizes with 1,000 sets of base stock levels representable enough for the entire population (i.e., 25,0000 sets of base stock levels). The box plot quartile ranges of the samples used to analyse the performance of the GA and local search approach are stated in Table E.5 and E.6.

Table E.5: Boxplot quartiles of a sample of 1,000 sets of base stock levels generated by the GA approach.Confidential

Table E.6: Boxplot quartiles of a sample of 1,000 sets of base stock levels generated by the LS approach. **Confidential**

F. Base stock levels of the best solutions

Table F.1: Base stock levels of the best (i.e., most cost-effective) solution generated by the GA approach. **Confidential**

Table F.2: Base stock levels of the best (i.e., most cost-effective) solution generated by the local search approach. **Confidential**